Advances and Challenges with Data Broadcasting in Wireless Mesh Networks

Chun Tung Chou, Junaid Qadir, and Joo Ghee Lim, University of New South Wales
Archan Misra, IBM T J Watson Research Center

This project is supported by Australian Research Council Discovery Grant DP0664791.

ABSTRACT

Wireless mesh networks have become a promising means to provide low-cost broadband access. Many WMN applications require broadcasting data (IPTV etc.) over the WMN. This article studies how efficient data broadcast, measured in terms of broadcast latency, can be realized by exploiting two features of WMNs: the use of multiple transmission rates at the link layer and the use of multiple radio interfaces on each node. We demonstrate that by exploiting these features, broadcast latency can be reduced severalfold compared to the current default practice in wireless LANs of using the lowest transmission rate for broadcast traffic. We also discuss two important insights we have gained from our investigation. First, we find that when multiple radio interfaces are used, a channel assignment algorithm designed for unicast traffic may often perform poorly for broadcast flows. Second, we find that the efficiency of a transmission rate for reducing broadcast latency can be reasonably predicted by the product of the transmission rate and its coverage area.

INTRODUCTION

Wireless mesh networks (WMNs) offer a promising low-cost broadband access alternative in many cities (e.g., New York), suburban communities, and campus environments, especially when built from commodity wireless cards and operating over unregulated spectrum. The static mesh nodes (mounted, e.g., on residential rooftops or light poles) form a multihop wireless overlay, with an individual mesh node also acting as an *access point* to mobile and consumer devices in its vicinity. Many of the applications (e.g., [1]) enabled by such mesh infrastructures are *broadcast* oriented and involve the point-to-multipoint transmission of multimedia (audio, video) data. Examples include broadcast IPTV, collaborative multiparty communications (e.g., IM, voice conferencing), local streaming (e.g., video feeds from neighborhood security or traffic cameras), and multiplayer online games.

In this article we describe some of our initial work in the Aiolos project [2], where we are studying WMN-specific opportunities and challenges for broadcast data transmission. (We use the term *broadcast data* to distinguish our work from earlier research on mobile ad hoc networks, where network-wide broadcasts were used largely for sporadic *control* traffic, generated by route establishment or repair operations of ad hoc routing protocols.) In particular, we show how efficient data broadcasting can and should leverage on the following two features of WMNs:

- *Rate diversity in link transmissions*: Most commodity wireless cards perform adaptive modulation, whereby the link transmission rate changes in response to the receiver signal-to-noise ratio. If the same transmission power is used for all link transmission rates, in general, the faster the transmission rate, the smaller the transmission range (although, as shown in [3], the rate-distance variation in real life is somewhat irregular).
- *Use of multiple channels and radios on individual mesh nodes*: Many research prototypes of WMNs equip each mesh node with multiple radio interfaces and tune these radios to orthogonal (or nonoverlapping) channels to reduce the overall interference experienced in the network. Recent results have shown that the increase in WMN capacity due to the resulting improvement in spatial reuse is usually nonlinear (e.g., [4] reported a sixfold increase in WMN capacity when each node has three interfaces).

For compactness, we shall henceforth refer to a WMN consisting of mesh nodes with single radio interfaces (operating on a common channel) as single-radio single-channel (SRSC), and a WMN consisting of mesh nodes with multiple interfaces as multiradio multichannel (MRMC).

Both rate diversity and multiple-interface node architectures have previously been investigated for *unicast* flows. In contrast, our focus is on broadcast (and, by extension, multicast) application traffic. In contrast with broadcast algorithms for wired networks, our well

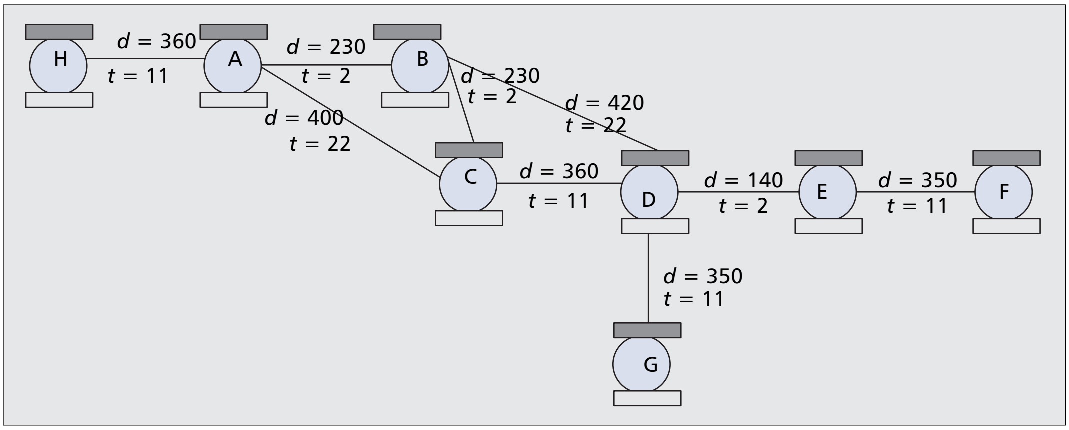

Figure 1. *Sample WMN topology illustrating rate diversity and multichannel mesh operation.*

Choosing latency as a performance measure implicitly rewards approaches that use the WBA to reduce the number of distinct transmissions, since this reduction directly translates into lower contention-induced delay.

designed broadcasting strategies in WMNs must exploit the wireless broadcast advantage (WBA), wherein a single transmission over the wireless medium can reach multiple downstream receiver nodes. Since many of our target broadcast scenarios involve interactive or streaming data, we measure efficiency in terms of *broadcast latency*, which we define as the maximum delay between the transmission of a packet by the source node and its eventual reception by all the destination (receiver) nodes. Choosing latency as a performance measure implicitly rewards approaches that use the WBA to reduce the number of distinct transmissions, since this reduction directly translates into lower contention-induced delay.

Broadly speaking, this article provides insights into the following fundamental questions:

- Is the efficiency of broadcasting in WMNs likely to be impacted by link rate diversity and multiradio, multichannel WMN nodes?
- How can we design broadcasting algorithms that exploit link diversity, multiple radio interfaces, and WBA to achieve low broadcast latency? What are the performance gains expected from such "smart" WMN broadcast algorithms?
- Can we derive some generic insight into the choice of link layer rates or channel assignment algorithms for broadcast traffic?
- What are the open issues and challenges associated with WMN broadcasting and multicasting?

The rest of this article answers these questions. We use simple examples to illustrate the novel degrees of freedom that rate diversity and multiradio operation offer for broadcast traffic. We detail the design of a rate-diversity-aware broadcasting algorithm in an SRSC WMN, and quantify the observed latency gains. We then introduce and study the performance of algorithms that additionally exploit the increased concurrency offered in MRMC architectures. We analyze the impact of rate diversity on multicasting and present a general guideline called the Rate Area Product principle. We list some of the open research challenges associated with efficient broadcasting in WMNs. Finally, we conclude the article with a discussion of our ongoing work and other open problems associated with data broadcasting.

Illustrating the Impact of Rate Diversity and Multiradio Nodes on Broadcast Latency

We first use a simple topology to illustrate the potential degrees of freedom that may be available in a WMN due to both the multirate nature of individual links and the availability of multiple radios on each node. Figure 1 shows an 802.11b-based topology consisting of eight nodes $\{A, B, C, D, E, F, G, H\}$, where A is the source of a network wide broadcast. Assume that each node has two radios, with the dark and light interface on each node denoting a radio tuned to channel C_1 and C_2, respectively. Each edge in Fig. 1 includes the distance d between the neighbors and the assumed packet *transmission time* on that link (derived from the Qualnet simulator, where transmission ranges for 1, 2, and 11 Mb/s are 283 m, 370 m, and 483 m, respectively). Note that the transmission times t in Fig. 1 are normalized to the transmission time for the fastest rate (11 Mb/s) and are inversely proportional to the link rate; thus, links with time $t = 2$, 11, and 22 have link rates of 11, 2, and 1 Mb/s, respectively. The interference range is 520 m: the reception of a packet by a receiver will be unsuccessful if there are additional active transmitters within 520 m of the receiver.

Note that we make two ideal assumptions here:

- The transmission time is computed based on the physical layer transmission rate.
- The medium access control (MAC) layer is ideal (no collisions or backoff).

The transmission time therefore ignores the overhead in packet headers, channel switching time, and contention resolution. Later, we shall demonstrate (via discrete-event simulation studies) that our fundamental insights hold even when we incorporate the overheads associated with a non-ideal MAC.

A Single-Channel WMN

To study the different broadcast options available in a single-channel WMN, we consider the case where only one radio (the solid interface tuned to channel C_1) is active in each node. In the first broadcast strategy called Alt_1 that does not exploit rate diversity, all transmissions would

The topology construction phase decides which node will be transmitting. The number of transmissions and the actual transmission rates used by a particular forwarding node is determined in the Multicast Grouping phase.

occur at the lowest rate (i.e., at 1 Mb/s). In this case the transmissions would occur as $A \to \{B, C, H\}$, followed by $C \to D$, $D \to \{E, G\}$ and $E \to F$. Since each of these transmissions takes t = 22 (i.e., at 1 Mb/s), the total broadcast latency is 88 time units.

In contrast, an alternate rate-diverse strategy called Alt_2 would have $A \to \{B, H\}$ (t = 11), followed by $B \to C$ (t = 2), $C \to D$(t = 11), $D \to \{E, G\}$ (t = 11), and finally, $E \to F$ (t = 11). It is easy to see that, by exploiting rate diversity, this schedule can reduce the broadcast latency to (11 + 2 + 11 + 11 + 11 =) **46** time units.

Even more interesting, let us consider another transmission scheme Alt_3, where *each node is allowed to transmit the same packet more than once but at different rates*. In this case, if A first transmits the packet only to B at 11 Mb/s (t = 2), this transmission would not be received by E (as it can only decode transmissions at the lower 2 Mb/s rate). Moreover, the transmission $A \to E$ could not proceed concurrent with the subsequent transmission $B \to C$ (t = 2) (as A's transmission would cause interference at C). However, the two transmissions $C \to D$(t = 11) and $A \to H$ (also t = 11) could proceed in parallel, and both of them can begin once C has received its packet. After D has received the packet, the remaining transmissions are $D \to \{E, G\}$ (t = 11) followed by $E \to F$ (t = 11). The broadcast latency (2 + 2 + 11 + 11 + 11 =) **37** time units.

The above examples illustrate two important and novel features that are unique to broadcast flows in SRSC WMN environments:

- Exploiting link rate diversity for link layer broadcasts can directly improve the broadcast latency.
- Allowing an individual node to transmit the same packet multiple times, at different rates to different subsets of downstream receivers, can potentially further reduce the broadcast latency.

EXPLOITING MULTIPLE INTERFACES

Continuing the example of the previous subsection, we now assume that each node has both of its radio interfaces (channels C_1 and C_2) active. In this case the nodes can exploit this additional level of concurrency to further reduce the number of interfering transmissions and the resulting latency. The transmission schedule Alt_4 is similar to Alt_3 above except that D would transmit on both radios but to different neighbors. The first four transmissions are (as in Alt_3): $A \to B$ (t = 2), $B \to C$ (t = 2), two simultaneous transmissions by $A \to H$ and $C \to D$ (t = 11). These transmissions can take place in either of the channels. When D receives the packet, it transmits to E and G on two separate interfaces. Without loss of generality, we assume that $D \to E$ (t = 2) is on channel C_1 and $D \to G$ (t = 11) is on channel C_2. Once E has received the packet, it would transmit to F on channel C_1 with t = 11. Note that the transmissions on C_1 and C_2 can take place simultaneously. Therefore, by exploiting multiple interfaces, the broadcast latency is further reduced to (2 + 2 + 11 + 2 + 11 =) **28** units.

LOW-LATENCY BROADCASTING IN A SINGLE-RADIO WMN

We now turn our attention to the SRSC WMN (where all nodes contend for a single common channel), and first develop an algorithm for low-latency broadcasting that exploits both WBA and link rate diversity. Developing such an algorithm will enable us to numerically evaluate the degree to which the incorporation of rate diversity may benefit broadcasting in real WMNs. Any broadcast algorithm has to compute the structure of the tree (i.e., which subset of nodes act as forwarding nodes), how many times the non-leaf nodes would transmit, and at what rates. Not surprisingly, the generic optimization formulation is shown to be NP-hard in [5], implying that we must develop heuristic strategies for any realistic WMN size. Given the hardness of the problem, we decompose the problem into three independent steps:

- *Topology construction*: This step computes a broadcast tree T of the given multirate WMN that exploits both the multirate nature of links and WBA. At the end of this step, we only determine who the transmitting nodes are (i.e., the non-leaf nodes) and the children/parent relations between different nodes, but not the precise rate used in individual transmissions.
- *Downstream multicast grouping*: This step takes the broadcast tree T from above, and determines the number of distinct link layer transmissions (because multiple transmissions by the same node may reduce broadcast latency) and the corresponding transmission rate for each non-leaf node in T.
- *Transmission scheduling*: For a theoretical study of the algorithm's performance benefits, this step assumes a centralized scheduler and schedules (determines the precise order of) the individual transmissions of T to avoid interference.

Clearly, this sequential decomposition of the low-latency broadcasting problem may be suboptimal, since the choices in one phase restrict the degree of freedom in subsequent steps. We shall see later that in realistic WMN topologies, the impact of *grouping* is fairly minor. Moreover, the centralized scheduler is used purely as an analytical tool to gauge the performance limits of rate-aware broadcasting; the relative performance of various algorithms remains unchanged when such a scheduler is replaced by a distributed MAC in real WMN deployments.

THE WCDS TOPOLOGY CONSTRUCTION ALGORITHM

One well-known heuristic approach to building a broadcast tree on a graph of nodes is based on the concept of a connected dominating set (CDS). For a graph $G = (V, E)$ (where V denotes the nodes and E denotes the links between the nodes), a CDS Z of G is a subset of V such that every element (node) of $V \backslash Z$ is in the neighborhood of at least one node in Z, and the set Z is connected. The CDS defines a broadcast tree where each node $v \in Z$ is a non-leaf node, with s being the broadcast source. To minimize the number of individual transmissions, researchers [6] have previously developed heuristics to approximately compute the minimum CDS (MCDS) (i.e., minimizing the number of elements in Z) when all the links have identical rates. For our multirate WMN environments, we need to extend this MCDS heuristic to take rate diversity into account.

Our heuristic is called the Weighted MCDS (WCDS) [5] algorithm, since it essentially tries to compute the MCDS while weighing each non-leaf node in Z by its link transmission rate. A central decision of the WCDS algorithm is to determine the non-leaf nodes of the broadcast tree. To achieve low broadcast latency, it is important to exploit WBA to cover as many nodes as possible while maintaining a preference for high transmission rates. The WCDS algorithm, which works in a greedy manner, constructs the broadcast tree incrementally from the source. In the following algorithmic description, we say that a node is covered if it has already received the packet. An outline of the WCDS algorithm is:

- Initialization: The source node receives the packet and is therefore covered. All the other nodes are yet to be covered.
- For all covered nodes n and possible transmission rates r that n can use, compute $f(n, r)$ = the product of r and the number of yet-to-be-covered nodes reachable by node n transmitting at rate r. The node-rate combination (n, r) that maximizes $f(n, r)$ will be selected as the next transmission. The nodes reachable by this selected transmission will become covered since they will have received the packet.
- If all the nodes in the network are covered, quit; otherwise, return to step 2.

An important design consideration in our algorithm is to capture the trade-off between the transmission range (which affects the number of nodes covered) and the transmission rate. Although a transmission using a higher rate takes a shorter time, it covers a smaller number of nodes and may increase the total number of transmissions needed to cover the network. The exact converse applies to lower-rate transmissions. Our algorithm captures this trade-off by using the function $f(n, r)$, which balances the rate and the number of nodes covered.

The topology construction phase decides which node will transmit. The number of transmissions and actual transmission rates used by a particular forwarding node are determined in the *multicast grouping* phase. This phase operates progressively from the bottom (leaf nodes) of the forwarding tree, running through all the possible transmission patterns to choose the one that results in the lowest latency for the underlying subtree. By now, all the transmissions and their rates have been determined. The *scheduling* phase decides on the order of transmission by scheduling the most critical transmission (according to estimated broadcast latency) first. For reasons of space, see [5] for details of the grouping and scheduling algorithms.

NUMERICAL RESULTS AND INSIGHTS

To study the performance benefits of WCDS, we performed simulation studies using random WMN topologies of different network sizes (number of nodes n) and network area A. Our WCDS heuristic was compared against two alternative schemes:

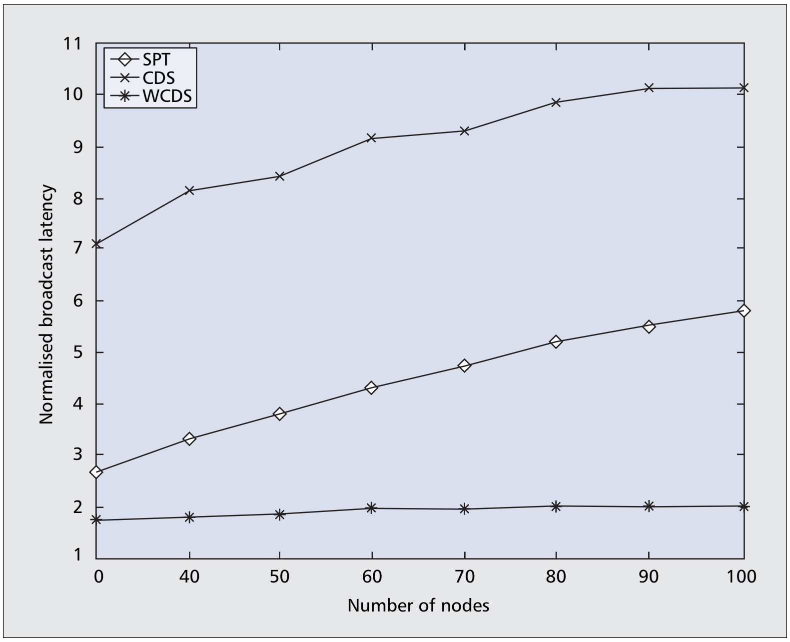

Figure 2. *The geometric mean of the normalized broadcast latency of WCDS, SPT and CDS using an ideal MAC layer.*

- *Algorithm SPT*: This heuristic uses Dijkstra's algorithm to compute the shortest path tree (SPT), and thus does not take WBA into account.
- *Algorithm CDS*: This heuristic assumes that all broadcasts are done at the lowest transmission rate (as is the common practice in current wireless LANs). The broadcast tree for CDS can be computed using the WCDS algorithm, assuming that only the lowest rate is available.

Figure 2 shows the comparative broadcast latency of the WCDS, CDS, and SPT algorithms as the number of nodes (n) is varied from 30 to 100 over an $A = 1.5$ km^2 area. (In these experiments we explicitly disable multicast grouping, thus allowing each node to transmit only once.) The delays are normalized by the *ideal SPT* delay (i.e., expressed as a ratio of the delay the SPT tree would incur if it operated in a wired network where each link is point-to-point) and represent the geometric mean of 100 simulation runs. The figure shows that CDS (which does not exploit rate diversity) and SPT (which does not exploit WBA) result in ~ 7–8 and ~ 4–6 times, respectively, the ideal latency. In contrast, by considering both rate diversity and WBA, WCDS can reduce the broadcast latency to ~ 2 times that of the corresponding wired network. Note that WCDS outperforms SPT here because WCDS exploits WBA and results in fewer transmissions than SPT; for details see [5].

We have also studied the potential impact multicast grouping might have on the broadcast latency in a WMN, and found this degree of freedom to be rarely used. For the WCDS algorithm, only 2 out of 100 topologies for a network area of 1 km^2 required multiple transmissions; moreover, the use of multiple transmissions resulted, on average, in only a 10 percent reduction in broadcast latency. Additionally, we also computed the

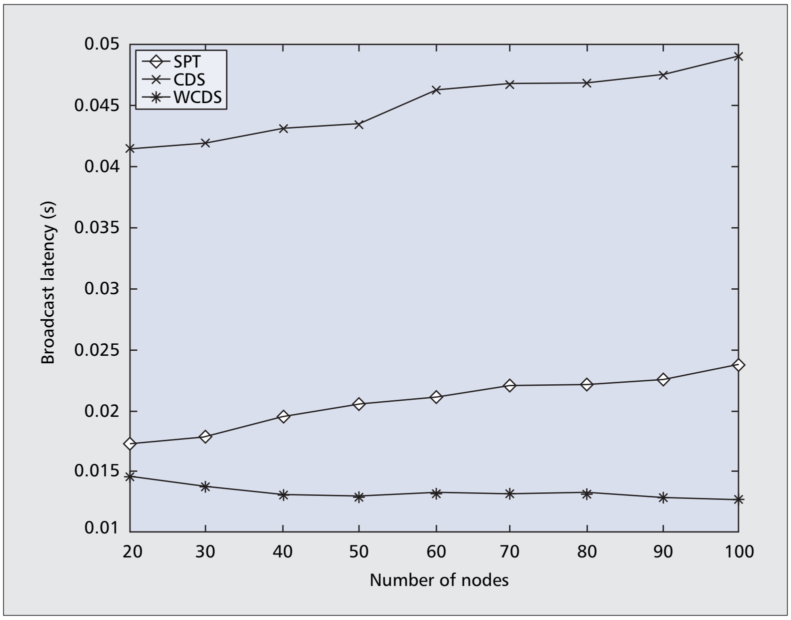

Figure 3. *The mean of the broadcast latency of WCDS, SPT, and CDS using a non-ideal MAC layer obtained from simulation.*

maximum data throughput that the resulting tree T could sustain, and found it to always be inversely correlated to the broadcast latency. Accordingly, an algorithm such as WCDS not only ensures low-latency packet delivery, but also maximizes the data capacity of the broadcast tree.

We have also studied the relative performance of our proposed algorithms for a non-ideal distributed MAC, using the Qualnet simulator to evaluate the broadcast latency in 802.11-based WMNs. Figure 3 plots the average broadcast latency in seconds over 100 randomly generated WMN topologies for three different algorithms: CDS, SPT, and WCDS. It shows that with a non-ideal MAC layer, WCDS still outperforms CDS (~3–4 times) and SPT (~1.2–1.8 times).

Overall, we can thus conclude that explicit incorporation of rate diversity (via an algorithm such as WCDS) can *reduce the latency of WMN broadcasting to about 1/3–1/4 of that incurred when links operate purely at the smallest rate (CDS)*. However, the degree of freedom associated with multiple transmissions per node (step 2) is less critical to performance, at least in sufficiently dense WMN deployments.

Reducing Latency in a Multiradio Multichannel WMN

We can expect the broadcast latency will be even further reduced in an MRMC network, where each node has multiple radio interfaces it may utilize simultaneously. Of course, the performance of the broadcast algorithm will also depend on the channel assignment strategy (which channels are assigned to the different radios on each node). We assume that channel assignment is done *independently* (before the broadcast routes are computed) and consider only *static* channel assignment schemes (where the channel assigned to a particular radio does not change at every packet transmission). As in the SRSC case, the broadcast routing algorithm can be decomposed (for tractability) into three sequential phases:

- The topology construction phase
- The grouping phase
- The scheduling phase

Once again, for reasons of space we concentrate on the *topology construction* phase.

While in an SRSC two nodes can communicate if they are within transmission range of each other, the situation in MRMC is different as it requires in addition that these two nodes also share a common channel. Since it is possible for two neighboring nodes to share a number of common channels, the abstract representation of an MRMC WMN is usually a *multigraph* $\mathcal{G}$ with multiple edges between the same pair of nodes when the node pair shares two or more channels.

Broadly speaking, the heuristic algorithms for such an MRMC mesh must utilize this potential availability of multiple interfaces on each node, and the fact that parallel transmissions (at possibly different rates and to different subsets of neighboring nodes) may proceed simultaneously.

The PAMT Algorithm

The Parallelized Approximate-Shortest Multiradio Multichannel Tree (PAMT) algorithm extends the WCDS algorithm to *adapt* to the number of radio interfaces and channels available. The PAMT algorithm is based on the observation that a node m covered by a transmission combination (n, r, c) (node n transmitting at rate r on channel c) may also be covered by a concurrent combination $(\hat{n}, \hat{r}, \hat{c})$ such that $r < \hat{r}$ and $c \neq \hat{c}$ (another node transmitting potentially at a faster rate on an orthogonal channel). In fact, such a transmission $(\hat{n}, \hat{r}, \hat{c})$ might actually be preferable if the total latency of reaching m via the path through $\hat{n}$ is smaller than the latency incurred via node n. The PAMT algorithm encapsulates this observation by considering a node reachable at a lower latency by an alternate combination $(\hat{n}, \hat{r}, \hat{c})$ to be already covered, even though it may not have received the packet. Like WCDS, PAMT works in a greedy manner, choosing the *best* transmission in each round that maximizes an efficiency measure $f(n, r, c)$: the product of the number of not-yet-covered nodes and the rate r. The key difference is that the $f(n, r, c)$ computation in PAMT does not consider those nodes that may receive the packet via an alternative path at lower latency. For algorithmic details see [7] where we designed and studied the performance of a number of broadcasting algorithms for MRMC networks.

Numerical Results and Insights

Figure 4 compares the latency of PAMT to two other algorithms detailed in [7]:

- *MSPT*: This heuristic is based on Dijkstra's shortest path algorithm, with the additional capability of choosing a channel between two nodes that have multiple common channels.
- *MWT*: This is the multichannel equivalent of the WCDS algorithm in that it greedily chooses a (node, rate, channel)-tuple, as opposed to WCDS, which chooses (node, rate)-tuples. Unlike PAMT, it does not consider the latency reduction possible from alternate paths.

As before, the latency bounds are computed normalized to Dijkstra's idealized bound for a corresponding wired network. From the figure we see that PAMT provides the best performance among the available heuristics. More significant, our results show that even with a small number of interfaces (Q, the number of interfaces, equals 3), the broadcast latency for PAMT is only 5 percent higher than the ideal case, compared to WCDS performance in Fig. 2, where the delay was almost twice the ideal lower bound. This result and additional simulation-based studies we have performed demonstrate *how the proper use of a small number of radios on a single mesh node can lead to a significant reduction in broadcast latency*. This reduction is due to a combination of both the reduced interference because of the presence of orthogonal channels and the exploitation of concurrent transmissions by an individual node to different subsets of downstream neighbors.

Impact of Channel Assignment Strategy

One of the interesting features of MRMC mesh networks is the dependence of broadcast performance on the underlying channel allocation strategy, since the channel allocation effectively determines both the *connectivity* and *interference* characteristics of the multigraph $\mathcal{G}$. For broadcast traffic, a channel assignment algorithm has to balance two conflicting objectives:

- Low interference so that individual transmissions have lower contention conflicts
- High connectivity so that a single transmission can reach a large number of neighbors sharing a common channel

To demonstrate this effect, Fig. 5 plots the normalized delay bounds of the MWT and PAMT algorithms for three commonly used channel allocation strategies:

- *CCA*: In the common channel assignment strategy, all the nodes are assigned a common set of channels (i.e., interface 1 is assigned channel C_1, interface 2 channel C_2, etc.).
- *VCA*: In the varying channel assignment strategy, channels are assigned to each interface at random, so different nodes and interfaces have different channels.
- INSTC: In the interference survivable topology control approach [9], channels are assigned to minimize the interference while maintaining connectivity in the induced multigraph.

For unicast traffic, INSTC has been shown [9] to offer the best performance among the three schemes. However, as Fig. 5 shows, for broadcast traffic, CCA generally results in the lowest latency, especially for the practical scenarios where the number of interfaces (Q) is relatively small. Overall, our studies thus lead to the important observation that *a channel assignment scheme designed for unicast traffic may often perform poorly for broadcast flows*.

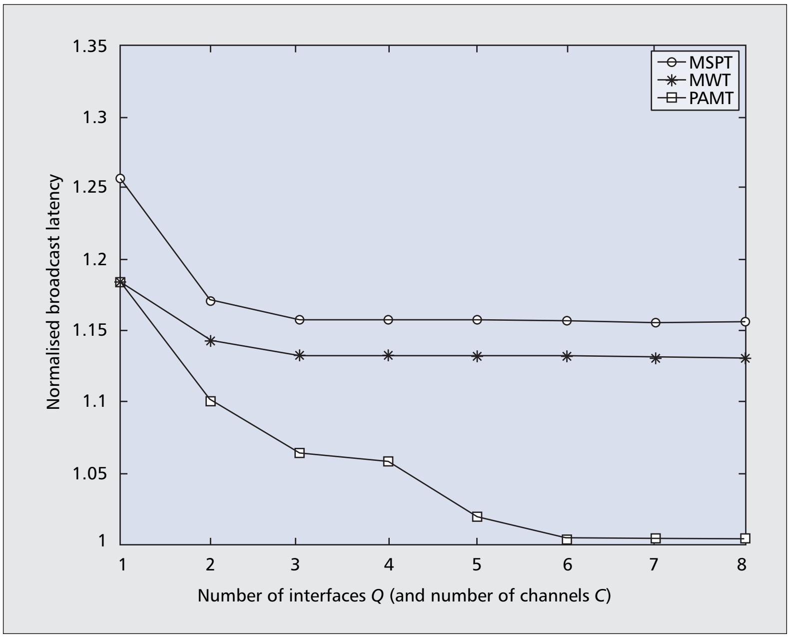

Figure 4. *Normalized broadcast latency against varying number of radio interfaces* Q (C = Q) *with N = 10 (A = 1.2 km²).*

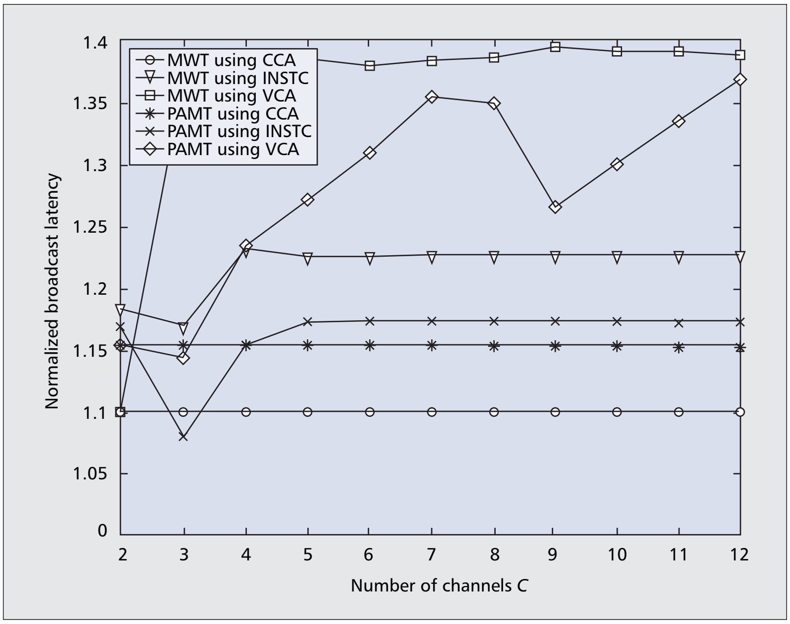

Figure 5. *The impact of channel assignment on MWT and PAMT algorithms for* Q = 2, N = *30 and A = 1.2 km²).*

Generalizing Our Insight: The RAP Principle

While the previous sections illustrate the benefits of rate diversity, we shall now shed light on a few fundamental questions related to the design of multirate WMNs. Specifically, given a WMN that can use k different transmission rates, it is useful to know: *Do we need to use all* k *rates? Are some rates more* efficient *than the others in reducing the broadcast latency? Is there a simple way to decide how efficient a transmission rate* r *is?* These are difficult question to address, since it is highly nontrivial to obtain a closed-form expression that relates the effect of link layer transmission rates on the minimum broadcast latency. To get around this problem, we used a special simulation setup in [5] to answer these questions. In this article we instead use an intuitive approach.

In this discussion we assume that the nodes of the WMN lies on a plane (the generalization

Transmission rate (Mb/s)	Transmission range (m)	RAP (Mb/s-km^2)
1	610	1.2
6	396	3.0
11	304	3.2
18	183	1.9
54	76	1.0

Table 1. *The transmission range and RAP of a commercial IEEE 802.11b/g product.*

to the 3D case is straightforward). For network-wide broadcast, where the sender's packet must eventually reach all the nodes, the broadcasting process for a planar network may be viewed, from a geometric standpoint, as a process of covering the physical area of the network with circles of different sizes that corresponds to the coverage area of each individual transmission. With this geometric picture in mind and the rate-range trade-off, we hypothesized that the efficiency of a transmission rate r *can be measured by the product of transmission rate and its transmission coverage area*. We refer to this efficiency measure as *rate-area product* (RAP). Essentially, RAP measures how fast the broadcast transmissions progress to cover the entire network. We have verified by simulation in [5] that a transmission rate with a higher RAP value is more efficient in reducing broadcast latency and vice versa.

By using the transmission range values specified for a commercial IEEE 802.11b/g product [10], Table 1 shows how RAP varies with transmission rates. It shows that the RAP value peaks at an intermediate transmission rate — both very low and very high transmission rates have lower RAP values. To understand whether this trend will remain fundamentally valid as wireless technologies evolve in the future, we also investigated the behavior of RAP for an ideal modulation scheme, where the transmission rate at a given distance is given by the Shannon capacity (the highest possible rate achievable on that link.) Specifically, if we assume that the radio power attenuates over a distance d according to $1/d^n$ where n is the path loss exponent (which typically takes a value of 2–6), the Shannon capacity at a distance d (for an additive white Gaussian noise channel) is given by

$$R = B\log_2\left(1+SNR\left(\frac{d_0}{d}\right)^n\right) \tag{1}$$

where R is the Shannon capacity, B is the channel bandwidth, and SNR is the signal-to-noise ratio at a reference distance d_0. By studying the behavior of RAP, as given by πRd^2, we find that RAP is indeed maximized at an intermediate transmission rate, which we call R_{max}. Interestingly, R_{max} is a function of the path loss exponent n only and is independent of the other parameters in the Shannon capacity equation. For $n = 4$ (which holds for 2-ray outdoor radio propagation model), R_{max} corresponds to a spectral efficiency of 2.3bps/Hz. This means that, as technology improves, giving rise to higher spectral efficiency, the RAP of higher transmission rates will eventually fall. While this analysis focused on the SRSC WMN, the RAP principle also appears to be valid in determining the efficiency of a particular link rate for MRMC networks. However, the case for MRMC WMNs is not as straightforward because broadcast latency is also influenced by channel assignment. For a more indepth discussion on this, see [8], where we studied how the choice of different link layer broadcast rates (for a WMN where all nodes are assumed to operate at an identical rate) affects broadcast latency.

While R_{max} maximizes the value of RAP, it is important to realize that R_{max} may not always be a valid rate for a specific WMN deployment, as the resulting transmission range may not be sufficient to guarantee network connectedness. It is thus necessary to always allow lower-rate (higher-range) broadcasts to ensure connectivity, especially in nonuniform or sparsely deployed WMNs. However, our analysis suggests that designers of future multihop wireless MAC protocols may safely limit the maximum link layer broadcast rate to the RAP maximizing rate, as higher transmission layer rates offer limited benefits.

Remark: The above discussion assumes the rate used in computing RAP is the physical transmission rate, and this holds when the overheads in the MAC layer are negligible. However, when these overheads cannot be ignored, the RAP principle still holds but we need to use an *effective rate* which takes the overheads into account.

Open Questions and Research Issues

Our goal in this article has not been the development of *practical protocols*, but rather the use of algorithmic techniques to establish some principles. In particular, we wanted to demonstrate the potential benefits broadcast traffic can reap from rate diversity and multichannel multiradio WMN architectures. A lot of additional research questions will need to be solved to provide effective practical support for point-to-multipoint traffic in WMNs. Some of these are:

- To begin with, the algorithms presented here aim to minimize broadcast latency for a single packet, and do not consider the traffic load of the flow. In reality, there will be multiple broadcast flows (originating at different sources) with different rates, and the algorithms will have to compute broadcast trees that factor in both the existing network load and the traffic load of the incoming flow. Two key research challenges for such capacity-aware routing are: *How do we accurately quantify the capacity of a node for broadcast traffic subject to contention from other flows? How do we built QoS-aware trees that effectively route around mesh traffic hotspots?*
- All our results presented here apply to broadcast traffic. In practice, the traffic may often be multicast, going to multiple, but not all,

mesh nodes. A specific challenge here is that all our current algorithms are sender-initiated (the tree grows outward from the sender), whereas multicast trees, by their very nature, are usually receiver-initiated. Thus, another open question is: *How do we extend our algorithms to exploit both rate diversity and WBA for the case of multicast traffic, so the resulting tree is computed and adjusted dynamically as receivers join or leave?*

• Our current algorithms are centralized, and essentially require knowledge of the entire network topology. While some degree of centralization may be possible in WMN environments (since mesh nodes are often static and the topology is not very dynamic), practical algorithms will require significant decentralization. Accordingly, an important question is: *How do we design distributed broadcast/multicast algorithms and protocols that exploit the unique characteristics of WMN environments?*

• Our studies have demonstrated the fact that existing unicast-oriented channel allocation schemes do not work uniformly well for both broadcast and unicast traffic. Since real WMNs will have a mixture of both traffic types, another important question is: *What channel allocation mechanisms for MRMC networks will prove to be less sensitive to variations in the relative proportion of broadcast and unicast flows?*

• Variable link quality (i.e., link loss rate) is another challenge in real WMN environments. While we have currently not considered the impact of link quality in our tree formation algorithms, the impact of lossy links is particularly important, since wireless standards (e.g., 802.11) usually do not provide retransmission-based link layer reliability for broadcast traffic. Thus, we need to analyze: *How do we modify the broadcast algorithms to account for the different, potentially time-varying quality of links to different neighbors?*

CONCLUSIONS

In this article we have demonstrated how the link rate diversity and multiradio architecture of individual mesh nodes are two features that deeply influence the efficiency of network-wide data broadcast in WMN environments. By combining these features with the wireless broadcast advantage, well designed broadcasting algorithms can lower the number of contending transmissions and dramatically reduce the broadcast latency. In single-radio mesh architectures, a heuristic such as WCDS, which tries to maximize both the number of child nodes of an individual transmission and the transmission rate, can reduce the latency almost sixfold. If individual nodes have even a relatively small multiple (say 3) number of radios, an algorithm such as PAMT can reduce the latency another 60–70 percent by exploiting the greater degree of transmission concurrency. Equally interesting is the fact that channel allocation strategies designed for unicast transmissions usually do not perform well for broadcast or multicast traffic.

We believe that work on efficient network layer broadcasting and multicasting in a WMN is still at a very early stage. Over the next few years, we believe that significant advances are needed to develop practical distributed broadcast/multicast routing protocols that not only improve overall network capacity but also prove *robust* in the face of dynamic link quality fluctuations that may be typical for many outdoor WMN deployments.

We believe that work on efficient network layer broadcasting and multicasting in a WMN is still at a very early stage. Over the next few years, we believe that significant advances are needed.

REFERENCES

[1] I. Akyildiz, X. Wang, and W. Wang, "Wireless Mesh Networks: A Survey," *Elsevier Comp. Networks*, 2005.

[2] "Efficient Distribution of Content in Multi-Rate, Multi-Channel Wireless Mesh Networks," The Aiolos Project, http://www.cse.unsw.edu.au/simaiolos/index.html

[3] J. Bicket *et al.*, "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *Proc. MOBICOM '05*.

[4] A. Raniwala, T. Chiueh, "Architecture and Algorithms for an IEEE 802.11-Based Multi-channel Wireless Mesh Network," *IEEE INFOCOM*, 2005.

[5] C. T. Chou, A. Misra, and J. Qadir, "Low Latency Broadcast in Multi-Rate Wireless Mesh Networks," *IEEE JSAC*, Special Issue on Multihop Wireless Mesh Networks, vol. 24, 2006, pp. 2081–91.

[6] H. Lim and C. Kim, "Flooding in Wireless Ad Hoc Networks," *Comp. Commun.*, vol. 24, no. 3–4, 2001, p. 353.

[7] J. Qadir, A. Misra and C. T. Chou, "Minimum Latency Broadcasting in Multi-Radio Multi-Channel Multi-Rate Wireless Meshes," *Proc. IEEE SECON*, Sept. 2006.

[8] J. Qadir, C. T. Chou, and A. Misra, "Exploiting Rate Diversity for Broadcasting in Wireless Mesh Networks," *IEEE LCN*, Nov. 2006.

[9] J. Tang, G. Xue and W. Zhang, "Interference-Aware Topology Control and QoS Routing in Multi-Channel Wireless Mesh Networks," *Proc. Mobihoc 2005*, May 2005.

[10] Cisco Aironet 802.11a/b/g Wireless LAN Client Adapters (CB21AG and PI21AG) Installation and Configuration Guide.

BIOGRAPHIES

CHUN TUNG CHOU (ctchou@cse.unsw.edu.au) is a senior lecturer and head of the Networks Research Group at the School of Computer Science and Engineering, University of New South Wales, Sydney, Australia. He received his B.A. in Engineering Science from the University of Oxford, United Kingdom, and his Ph.D. in control engineering from the University of Cambridge, United Kingdom. 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 (archan@us.ibm.com) is a research staff member at the IBM T. J. Watson Research Center, Hawthorne, New York, where he has been working for the past seven years on pervasive and presence-based network infrastructures. His current research projects include streaming event middleware for SIP-based "rich presence" in converged networks, wireless mesh networking protocols, and remote healthcare monitoring. He received his Ph.D. in electrical and computer engineering from the University of Maryland at College Park in May 2000. He has published extensively and is a co-author of papers that received Best Paper awards at ACM WOWMOM 2002 and IEEE MILCOM 2001. He is an editor of *IEEE Wireless Communications* and currently chairs the IEEE Computer Society's Technical Committee on Computer Communications.

JUNAID QADIR (junaidq@cse.unsw.edu.au) is currently a Ph.D. candidate at the School of Computer Science and Engineering, University of New South Wales, Australia. He completed his B.E. in electrical engineering from the University of Engineering and Technology, Lahore, Pakistan, in 2000, and his M.S. in computer eengineering from the National University of Science and Technology (NUST), Pakistan, in 2003. His broad areas of research interest include routing, network optimization, and QoS provisioning in multihop wireless networks.

JOO GHEE LIM (jool@cse.unsw.edu.au) is a Ph.D. candidate with the School of Computer Science and Engineering, University of New South Wales, Australia. He received his B.Eng. and M.Eng. in electrical and computer engineering from the National University of Singapore. Subsequently, he worked as an associate scientist at the Institute for Infocomm Research, Singapore. His current research interests include MAC, QoS, and internetwork coexistence issues in WLANs and wireless multihop networks.