

Capacity Scaling with Multiple Radios and Multiple Channels in Wireless Mesh Networks

Sumit Roy, Arindam K. Das, Rajiv Vijayakumar, Hamed M. K. Alazemi, Hui Ma and Eman Alotaibi

Abstract—Many portable client devices such as PDAs, laptops and cell phones are already equipped with multiple wireless radios, as are infrastructure side elements like access points and base stations. In this article, we argue that the increasing availability of such multiple radio nodes, in conjunction with a suitably structured multi-hop or mesh architecture, has the potential to mitigate some of the key limitations of present day wireless access networks that *do not* exploit the presence of multi-radios (*i.e.*, network nodes that use single radios). While in this work we concentrate on emerging next-generation 802.11 WLAN based network devices and architectural concepts, the ideas are also pertinent to other unlicensed wireless networks, such as those based on the 802.16 standard.

I. INTRODUCTION

The emergence of cost-effective wireless access networking technologies has changed mobile communications and computing in significant ways. These technologies include 802.11 as well as other incipient approaches such as WiMAX, based on the 802.16 standard, and Wireless USB2 type connectivity, based on 802.15.3a Ultra Wideband systems. The success of these systems to date has largely been in deployments in the home and small enterprise segments where coverage is limited and few users are served simultaneously, *i.e.*, *the network size is small*. There is now considerable interest in expanding the use of these technologies to so-called “dense networking” scenarios such as large enterprises and public hotspots to serve more users over a wider area. But are these technologies - in some suitably evolved version that exploit new architectural principles - the answer to the challenge of providing ubiquitous last/first mile access to an ever growing number of users? Stated differently, what will it require for such networks to scale? These questions, crucial to further deployments of cost-effective wireless networks, require satisfactory new design solutions.

In this work, we focus primarily on issues relating to the scalability of 802.11 WLAN networks as there is already good evidence that the current dominant model of user access based on the *infrastructure mode* of 802.11, is poorly suited

S. Roy, E. Alotaibi and Hui Ma are with the Department of Electrical Engineering, University of Washington, Box 352500, Seattle, WA 98195. *e-mails*: {roy, eman76, mahui}@ee.washington.edu.

A.K. Das is a Post Doctoral Research Associate with the Departments of Electrical Engineering and Aeronautics and Astronautics, University of Washington, Box 352500, Seattle, WA 98195. *e-mail*: arindam@ee.washington.edu.

R. Vijayakumar is a Post Doctoral Research Associate at the Department of Electrical Engineering, University of Washington, Box 352500, Seattle, WA 98195. *e-mail*: rajiv@ee.washington.edu.

H. Alazemi is an Assistant Professor, Department of Computer Engineering at Kuwait University. *e-mail*: hamed@eng.kuniv.edu.kw.

This work was partially supported by a gift from Intel Corpn.

for dense deployments and network scalability. *We emphasize that this is, in part, due to both architectural reasons as well as inherent limitations of current protocol stack design.* To understand the architectural limitations, it is useful to recall that today’s WLAN networks were designed as outgrowths of the wired Internet - *i.e.*, for providing last hop connectivity to the mobile end user - and were not intended to provide expanded coverage and become intermediate access networks on their own. Accordingly, in today’s 802.11 deployments, all access points (APs) are directly wired to the backbone network (*i.e.*, they are IP-addressable hosts on the Internet) but have no direct inter-AP connectivity. The type of traffic that is currently supported on such networks is of two types:

- downstream traffic to a mobile end-user involving a single final wireless hop, and
- an AP-based relay mechanism for all ‘local’ traffic, *i.e.*, relaying data between mobile clients that are associated with a common AP.

In an enterprise scenario with many users, supporting data between and to users over a wide area is likely to be of the wireless *multi-hop* variety. Thus, downstream traffic to a mobile from a gateway AP would pass through multiple ‘router APs’. Further, traffic between two mobile end users that are associated with *different APs* would also be routed in a multi-hop manner¹.

The protocol limitations in current 802.11 networks arise from the fact that all communication between wireless nodes in the same “cell” or 1-hop neighborhood occurs via a contention based mechanism which is governed by the base 802.11 Multiple Access (MAC) protocol, the Distributed Coordination Function (DCF) or CSMA/CA. Clearly, CSMA/CA within a cell presumes that there is no hidden terminal problem, *i.e.*, there is no mutual interference from co-channel users in other cells. This design presumption underlying current WLAN networks (a set of non-interacting cells) is increasingly invalid as growing user density leads to increasing co-channel or multi-access interference (MAI).² Since MAI impacts both the aggregate 1-hop throughput (relevant for last-hop wireless access for a large number of simultaneous users) as well as end-to-end user throughput (relevant for flow-type data traffic), integrated design of the link, MAC and network (routing) layers must be considered along with necessary supporting architectural changes. *Thus the fundamental issue behind*

¹All such traffic must currently get routed through the wired backbone due to lack of support for direct inter-AP communication, even if the two APs are in close physical proximity.

²This is due to inter-cell interference between nearby cells that are on the same channel.

successful scaling of such networks is to appropriately manage MAI jointly through collaborative design of all relevant layers in the protocol stack.

It is well-known from cellular systems engineering that the key to one-hop capacity scaling is *enhancing spatial reuse*, *i.e.*, reducing the re-use distance between co-channel users as much as possible. In narrowband systems such as FDMA, the extent of spatial reuse is directly proportional to the number of orthogonal channels available. Currently, only a very limited number of such orthogonal channels are available: 3 in 802.11b in the 2.4 GHz band and between 9 and 12 in 802.11a in the 5 GHz band depending on available bandwidth and channelization. Although greater worldwide allocation is anticipated for unlicensed use in the future, it is clear that relying primarily on increased bandwidth availability for capacity scaling is not a feasible option. Accordingly, for any given system bandwidth, optimizing the network performance necessarily requires improving the entire protocol stack, and efficient reuse of the available channels. When compared with mobile cellular systems, the ad-hoc nature of mesh node deployments leads to a higher degree of spatial variability of the MAI and significant location-dependent node throughput. This is where mesh architectures provide an advantage by allowing for more *fine-grained and dynamic* interference management and topology control via techniques such as node clustering and power control.

II. AP MESH ARCHITECTURES

The increasing availability of multi-mode radios, such as integrated 802.11b/g/a cards, in client and infrastructure devices will enable new mesh architectures. For example, Tier-1 client-AP connectivity may use the 802.11b/g radio while the Tier-2 backhaul AP mesh can use the 802.11a radio, thereby separating the different kinds of traffic (client-AP vs. inter-AP) and simultaneously utilizing the potential of multi-band radios. We first briefly discuss the performance of such a *single-radio* mesh (one radio per node) as a prelude to showcasing the advantages of a *multi-radio* mesh (multiple radios per node). The nodes in a Tier-2 mesh backhaul or access network consists of two types of nodes as shown in Figure 1 - a predominant lightweight subset (pure mesh points) whose only function is to route packets wirelessly to neighboring nodes and another subset of mesh AP nodes that allow direct client connectivity. A small fraction of these mesh AP nodes will be connected to the wired backbone and serve as gateways for traffic ingress/egress.

III. SINGLE RADIO MULTI-HOP MESH NETWORKS: 1-HOP AGGREGATE CAPACITY SCALING

Single radio multi-hop wireless networks are not new - in fact, they have been studied since the '70s under the nomenclature of packet radio networks. The end-to-end throughput in such single radio networks reduces with the number of hops. The primary reason for this is that a *single radio wireless transceiver operates in a half-duplex mode*; *i.e.*, it cannot transmit and receive data simultaneously and an incoming frame must be received fully before the node switches from




Fig. 1. Two types of mesh nodes: APs and Mesh Points

receive to transmit mode. Hence, a simple calculation for a linear chain with n half-duplex hops suggests that the end-to-end throughput will (at best) be inversely proportional to n .

Enhancing end-to-end throughput is related to increasing 1-hop aggregate throughput, which in turn depends critically on the extent of spatial reuse (*i.e.* the number of simultaneous transmissions per channel) that can be achieved in a given network area. Clearly, achieving a minimum separation distance between simultaneous co-channel transmitters on average would lead to maximum aggregate 1-hop throughput. This depends on the network topology and various characteristics of layers 1-3, namely the type of radio, signal quality requirements at receiver and signal propagation environment (layer 1 attributes), MAC attributes for interference management (layer 2 attributes) and choice of the routing metric for path determination (layer 3 attributes), suggesting that optimizing it requires a *multi-dimensional, cross-layer* approach. In this work, we will be content with highlighting the impact of only a few key aspects due to space limitations. In particular, we investigate the role of physical carrier sensing (PCS) in the IEEE 802.11 MAC protocol which is used by nodes to determine if the shared medium is available before transmitting to ensure that only acceptable interference occurs to ongoing transmissions. A node transmits only if the net signal power at its receiver is below a pre-set carrier sensing threshold. The choice of an optimal carrier sensing threshold depends on various (local) network properties; since these are often unknown a-priori, the thresholds should be tuned on-line in practice based on available information about local network conditions.

A. Single Radio, Single Channel Mesh Networks: A Baseline

Recent studies such as [1] show that the per-node share of the aggregate throughput of a single-channel multi-hop network of 802.11 nodes typically; behaves as $\frac{1}{n^\alpha}$, where n is the number of nodes, and the exponent α is influenced by topology and traffic characteristics. For example, an upper bound for large networks (*i.e.* via an asymptotic analysis) derived in [2] suggests that $\alpha = 0.5$ for a purely ad-hoc topology and random choice of source-destination pairs. Further insight can be obtained for finite networks with special

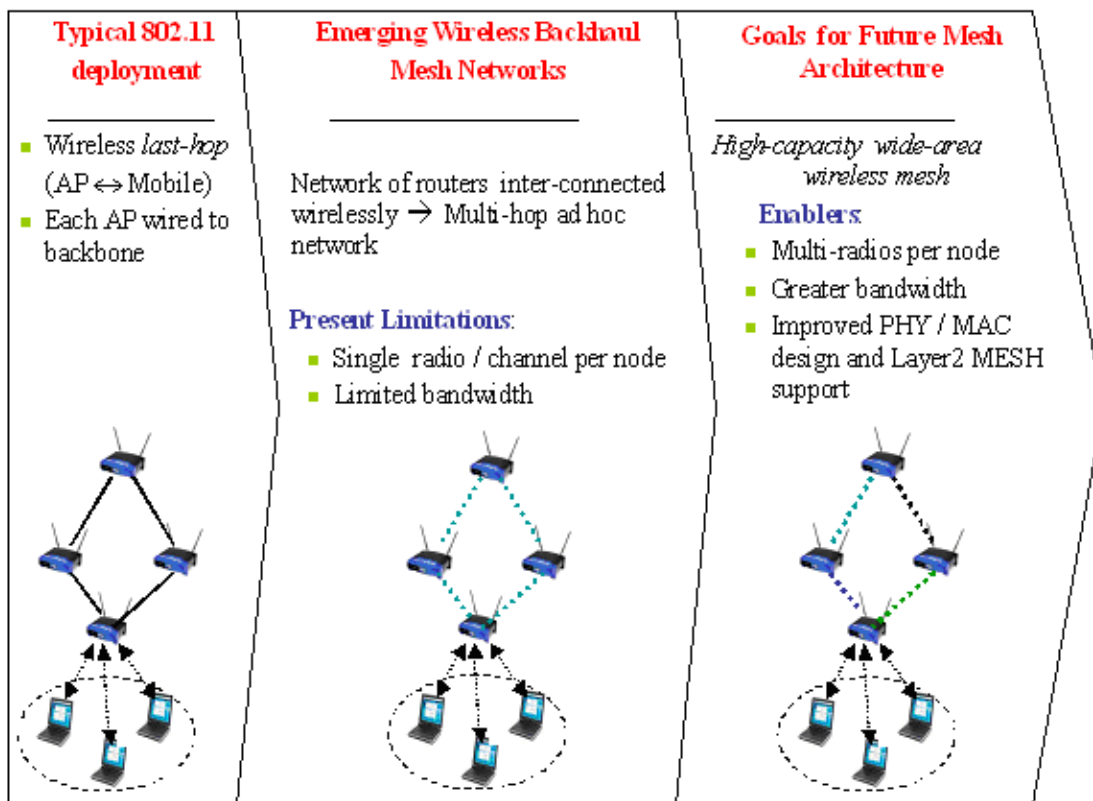


Fig. 2. Evolution of 802.11 AP mesh architecture.

topologies like a single-channel n node linear chain, for which the per-node throughput is $O(\frac{1}{n})$,³ implying $\alpha = 1$, since only a single transmission can occur at any time. This trend has been borne out via simulation results [3] with a 802.11 MAC.

We comment that the above results are founded on an important and *pessimistic* assumption central to scaling: that *all nodes in the network interfere with each other*, and that any pair of nodes (irrespective of their separation) communicate with equal probability. Typically, this is true only in small networks; in larger networks, traffic is more ‘localized’ (*i.e.*, nearby nodes communicate much more frequently). This implies that spatial reuse of channels is possible, leading to enhanced aggregate throughput. The role of spatial reuse in enhancing aggregate network throughput, facilitated by multiple (orthogonal) channels as well as multiple radios per node is discussed next.

B. Single Radio, Multi-channel Mesh Networks

Ideally, any end-to-end path in a multihop network should utilize all the available orthogonal channels⁴ (say C) in a

³This is true for a ‘small’ chain, or equivalently, a large carrier sensing range that prevents any spatial reuse of the single channel. Under the same assumptions, a chain with C channels will achieve an aggregate throughput of $O(C)$ which also does not scale with the number of nodes n .

⁴While very limited spatial reuse can be achieved even with $C = 1$, meaningful network scaling is only possible with increasing C .

manner that maximizes spatial reuse, *i.e.*, maximizes the number of simultaneous transmissions in the network area. Unfortunately, a key limitation of commodity single-radio wireless devices is that they operate in *half-duplex* mode, and therefore cannot transmit and receive simultaneously even if multiple non-interfering channels are available. A possible (but naive) approach to multi-hop route formation is for all nodes to use the same channel, even if multiple channels are available, at the cost of sacrificing spatial reuse. This approach does however avoid the serious drawback of *poor end-to-end delay* when adjacent node pairs use *different* channels to communicate. This necessitates channel scanning, selection and switching the radio at each node; this delay (per node) grows with C . For example, the switching delay for present 802.11 hardware ranges from a few milliseconds to a few hundred microseconds [5]. Further, the impact of frequent switching may be viewed as effective *route lengthening* because the switching delay manifests itself as *virtual hops* along the route [6]. On the other hand, exploiting the multiple orthogonal channels clearly enhances aggregate 1-hop throughput vis-a-vis the single channel scenario but at the cost of enhancing the end-to-end delay.

For all the above reasons, *multi-radio* meshes, which introduce several new degrees of freedom that fundamentally address the key limitation of commodity single-radio wireless devices, are expected to be a key component to achieving both network scalability and adaptivity in practice (as in software




Fig. 3. An example motivating the improvement in throughput that can be obtained with multiple radios and/or multiple channels.

(a) With one radio at node 2, each of the two flows, $1 \rightarrow 2 \rightarrow 3$ and $4 \rightarrow 2 \rightarrow 5$, receive an end-to-end throughput of $R/2$ bps (where R is the source rate) if they are scheduled at different times. However, if the two flows are simultaneous, the receive rate for both flows drops to $R/4$ bps. With two radios and availability of two orthogonal channels, the receive rate for both flows increases to $R/2$ bps, the same as each flow would have received if they were scheduled at different times.

(b) An illustration of a scenario when having multiple orthogonal channels is helpful even with one radio. For example, if two channels are available, one each can be used for the two transmissions. The receive throughput for each flow in this case is $R/2$ bps.

defined radios) for future wireless networks.

IV. MULTI-RADIO MESH

Multiple radio nodes are effectively full duplex; i.e., they can receive on channel c_1 on one interface while simultaneously transmitting on channel c_2 on the other interface, thereby immediately doubling the node throughput. As an example, consider the path $1 \rightarrow 2 \rightarrow 3$ in Figure 3. Let R denote the maximum possible transmit rate over one hop (i.e. from $1 \rightarrow 2$). With one radio, node 2 spends roughly half the time receiving from node 1 and the other half transmitting to node 3. Consequently, if the source (node 1) rate is R bps, the average receive rate at node 3 is approximately $R/2$ bps. With 2 radios at node 2 and 2 orthogonal channels, radio 1 can be tuned to channel 1 and radio 2 can be tuned to channel 2, in which case the receive rate at node 3 will be theoretically equal to R bps. Now, consider the case when there is a concurrent transmission on the route $4 \rightarrow 2 \rightarrow 5$. In this case, node 2 has to spend a quarter of its time receiving from nodes 1 and 4 and transmitting to nodes 3 and 5. The average receive rate at nodes 3 and 5 in this case is $R/4$ bps. Again, having multiple non-overlapping channels does not help in this specific scenario since the limiting factor is the availability of only one radio at node 2. Finally, consider the case when node 2 is equipped with two radios and there are 2 available orthogonal channels. In this case, radios 1 and 2 can be tuned to channels 1 and 2 respectively. If radios 1 / 2 are used on a half-duplex mode to support the routes $1 \rightarrow 2 \rightarrow 3$ / $4 \rightarrow 2 \rightarrow 5$ respectively, the average receiver throughput for each flow doubles to $R/2$ bps, the same as each flow would have received if they were scheduled at different times.

A key benefit of using multiple radios with multiple orthogonal channels is that a proper assignment of channels can be used to reduce the average (or maximum) collision domain size

of all transmission links. For example, consider the (partial) channel allocation in Figure 4(c) when there are two available channels. Figure 4(d) shows the total interference set of 16 edges which must remain silent when the link $a \leftrightarrow b$ is active on channel 1, assuming that all nodes have one radio. Of these 16 edges, only those which are assigned channel 1 (it can be verified that there are 12 of them) constitute the co-channel interference set. The remaining 4 edges which are assigned channel 2 need to remain silent, not due to interference issues, but because there is only one radio on nodes a and b . In other words, if more radios were available at nodes a and b , it should be possible for these 4 links to be active simultaneously. In a multi-radio, multi-channel framework, the total collision set therefore captures the effect of co-channel interference as well as hardware (radio) limitations. If two radios are available at a and b , any one of the two links incident on node a (or b) and assigned channel 2 can also be active. Figure 4(e) shows one possible configuration of the total interference set of 14 edges with two radios, twelve of which constitute the co-channel interference set (those assigned channel 1). Finally, Figure 4(f) shows the total collision set of 12 edges when three radios are available at a and b . Notice that, in this case, the total collision set is equal to the co-channel interference set.

Clearly, the allocation of channels to interfaces/radios will greatly influence end-to-end throughput, as will the choice of the metric for route formation - we discuss this issue subsequently. For example, for a node with R radios, one radio can be dedicated to perform channel scanning necessary for the channel-to-link allocation, thereby eliminating this increasingly significant overhead component (that scales with C) from the latency budgets for all the other radios. In summary, we suggest that with proper design, the *performance of multi-radio mesh scales as the size of the network increases by suitable design of Layers 1-3.*

A. Radio/Channel assignment and routing

A primary contributor to inefficiency (i.e., lower end-to-end aggregate throughput) in single-radio multi-hop networks is the significant overhead from standard routing protocols developed for wired networks. It is well known that the existing wired routing protocols, i.e., both proactive (table-driven) and reactive (on-demand) generate an amount of overhead that increases with the size of the network since they typically involve ‘all’ nodes in the route formation and rely on some form of intelligent broadcast for disseminating pertinent information for route computation. Thus for larger wireless mesh networks, the routing overhead information will ultimately consume most of the available bandwidth and consequently diminish the throughput, rendering these impractical [7]. Therefore, an adequate distributed solution based on local information only would be most desirable for multi-hop wireless networks.

Since wireless is essentially a broadcast medium, any transmission between two neighboring nodes impacts (in principle) transmissions anywhere else in the network. This has some immediate repercussions on choice of appropriate routing metrics in wireless and how they should differ from those in wired

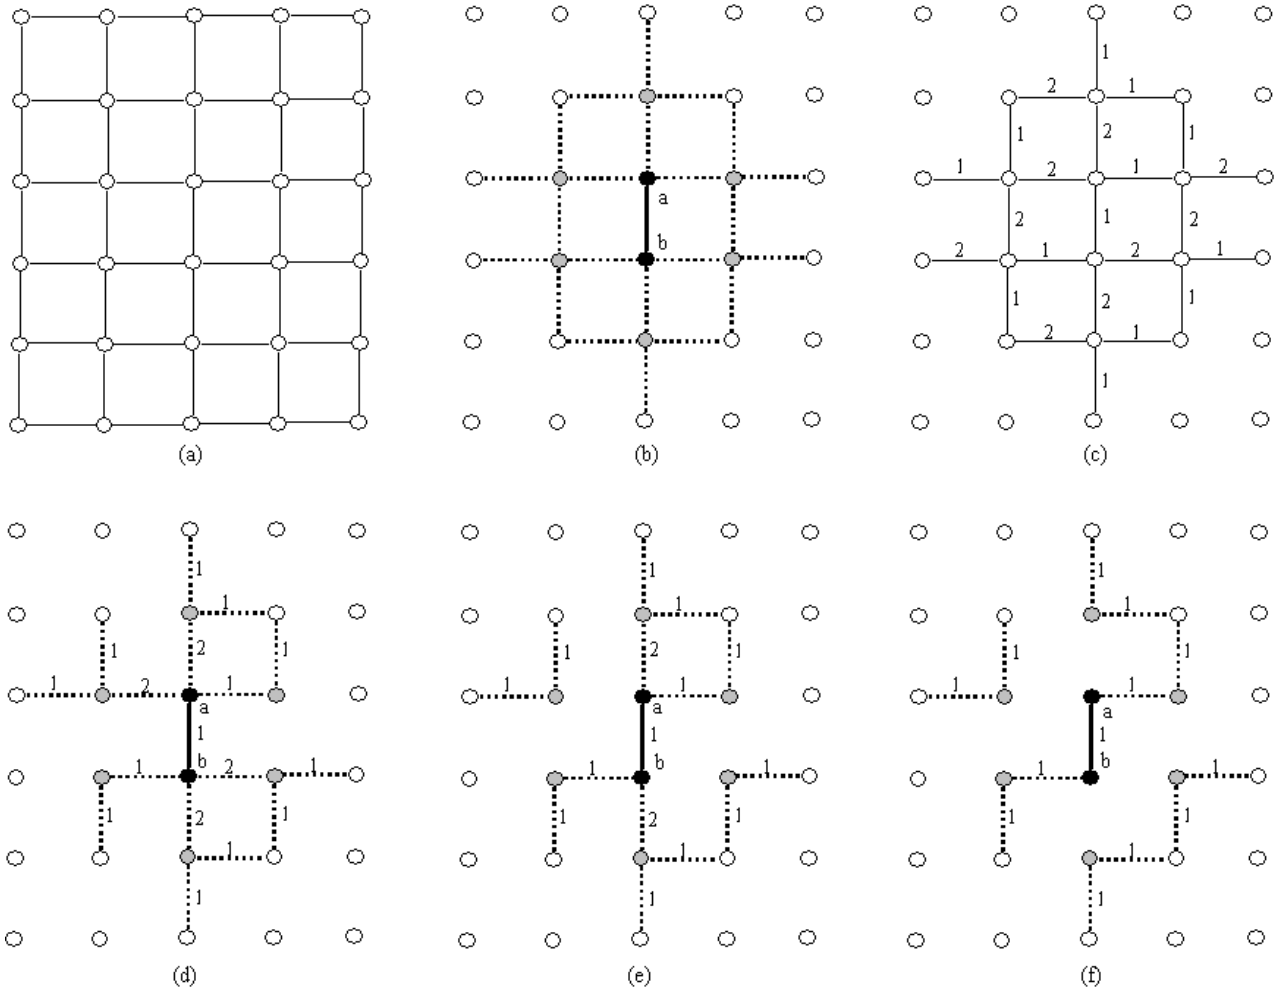


Fig. 4. (a) All possible transmission links in a 5×6 grid. (b) Set of 22 interfering links (shown dotted) for the edge $a \leftrightarrow b$, assuming that there is only one available channel. The lightly shaded nodes are neighbors of either node a or node b and must remain silent when the link $a \leftrightarrow b$ is active. (c) A partial channel allocation with 2 data channels. (d) Set of 16 edges which should remain silent when $a \leftrightarrow b$ is active on channel 1, if all nodes have one radio. 14 of these edges (those assigned channel 1) constitute the co-channel interference set. The rest of the edges (those assigned channel 2) are forced to remain silent due to insufficient number of radios on nodes a and b . (e) One possible configuration of 14 interfering edges which should remain silent when $a \leftrightarrow b$ is active on channel 1, if a and b each have two radios. 12 of these edges (those assigned channel 1) constitute the co-channel interference set. The rest of the edges (those assigned channel 2) are forced to remain silent due to insufficient number of radios on nodes a and b . (f) The set of 12 interfering edges (all due to co-channel interference) which should remain silent when $a \leftrightarrow b$ is active on channel 1, when a and b each have three radios.

networks. The drawbacks of classical shortest path routing algorithms for wireless networks have been well-documented; see *e.g.* [8]. Such algorithms simply select the path with the fewest hops without regard to the available link bandwidth, which can vary significantly depending on the interference environment at the local receiver. It is simple to construct examples where paths with larger number of hops can provide shorter end-to-end delay, depending on the residual bandwidths available on the links in the respective routes. A key to optimal route formation lies in exploiting *channel diversity* in multi-radio mesh networks, as indicated by channel usage along a path. In other words, longer paths (measured in number of hops) that reuse the available channels for better co-channel interference management may provide improved end-to-end

throughput/delay than shorter paths which use a fewer number of hops. In summary, desirable routing metrics for wireless must be *channel-aware*, *i.e.*, dependant on the underlying channel allocation. Unlike a wired network where each hop is assumed to be isolated from simultaneous transmissions on other links, the interaction between the link and upper layers is a vital and unavoidable element in wireless multi-hop routing.

The introduction of multiple radios adds a new degree of freedom to cross-layer design since there will be fewer co-channel transmissions. There have been several initial studies along this direction that are beginning to expose various aspects of this multi-faceted problem. Optimal joint channel assignment and routing is (unsurprisingly) NP-hard and various algorithmic heuristics are being proposed. In particular,

Raniwala *et al.* [5], [9] demonstrate that the channel assignment should depend on the load of each virtual link, which in turns depends on the routing metric. This dependency is clearly shown in their proposed centralized load-aware joint channel assignment and routing algorithm, which is constructed via a multiple spanning tree-based load balancing routing algorithm that can adapt to traffic load dynamically. Kyasanur *et al.* [6] study the multi-radio mesh network under the assumption that each node has the ability to switch some of its interfaces dynamically (to communicate with a nearest neighbor, typically) while the rest of the interfaces are fixed (*i.e.*, assigned the same channel for a relatively prolonged duration). They present a distributed interface assignment strategy that includes the cost of interface switching but is independent of traffic characteristics. Their routing strategy selects routes that have low switching and diversity cost and takes into account the global resource usage. Since all data is received on the fixed interface(s), neighboring nodes can communicate without any specialized coordination algorithm which would otherwise have been necessary to select a common channel for communication. Additionally, because of the fixed receive channel policy, no synchronization for channel switching is needed.

B. Channel-Aware Path Metrics for Routing

As we have seen, a good channel-aware path metric should incorporate notions of (i) total link cost (*e.g.*, sum of the transmission delays along the links in a path) and (ii) *path channel diversity*, which is a critical element in managing ‘self-interference’ in the network. The key challenge lies in suitably capturing this self-interference between hops on the same path using the same channel and finding a balance between these two components such that low cost links are not overused in any route. A good example of such a channel-aware metric is the *Weighted Cumulative Expected Transmission Time* (WCETT) path metric suggested by Draves *et al* [7]. Given an n hop path between a source and a destination, P_n , the WCETT of the path, $WCETT(P_n)$, is defined as:

$$WCETT(P_n) = (1 - \beta) \sum_{i=1}^n ETT_i + \beta \max_{1 \leq c \leq C} X_c$$

where the first summation term on the r.h.s represents the total link cost, the ‘max’ term on the r.h.s represents the bottleneck channel diversity along the path P_n , $0 < \beta < 1$ is a weighting factor on the two cost components and C is the number of orthogonal channels. The parameter ETT_i is the expected transmission time for the i^{th} hop on the path and is given by:

$$ETT_i = ETX_i \times \frac{S}{B}$$

where S is the packet size (in bits), B is the bandwidth (in bps) and ETX_i is the expected number of retransmissions (due to errors and losses) per packet on the i^{th} hop of P_n . For a certain channel c , the variable X_c represents the sum of the ETT’s for those hops which use channel c along the path P_n .

More formally,

$$X_c = \sum_{\text{hop } i \text{ is on channel } c} ETT_i; 1 \leq c \leq C$$

However, it suffices to state that, at this time, the research into choice of good path metrics and routing algorithm is in early stages and much remains to be done. For example, while the WCETT metric appears adequate for small networks, its performance degrades with increasing network size, indicating that improved channel diversity measures are desirable. The performance degradation with network size is primarily due to the fact that the WCETT metric does not consider ‘allowable spatial reuse’ along a path. In other words, for paths with a relatively high number of hops in a sufficiently spread out network area, there may be link segments which could successfully reuse a channel without undue interference. However, the WCETT metric, as proposed, does not allow for such spatial reuse in its channel diversity component and all hops sharing a channel are penalized equally, irrespective of their physical separation.

In the next section, we present some preliminary OPNET-based simulation results which highlight the impact of network topology and channel assignment on the performance of multi-radio, multi-channel wireless mesh networks.

V. SIMULATION RESULTS

While earlier we had emphasized the possible need for jointly optimizing channel allocation and routing algorithm, in practice it is likely that these will be optimized separately for various reasons (algorithmic simplicity being one, at the cost of some sub-optimality), at least initially. In this section, we opt to focus on the impact of network scale, carrier sense threshold and channel assignment when a ‘standard’ shortest path based routing algorithm is employed at the IP layer. This implies that channel assignment and optimal route determination are effectively decoupled. We plan to investigate the impact of the choice of the routing metric on network performance in the near future. All simulations were conducted using the OPNET network simulator.

A. One-hop Throughput Scaling as a function of Network Size

In this section, we present some preliminary results on the impact of network size on the aggregate 1-hop throughput as a function of the carrier sense (CS) range. The implication of CS range is that any source within this range of the reference transmitter will sense the ongoing transmission and defer its own. We concentrate on the single radio single channel case; similar experiments are currently being run for multi-radio, multi-channel meshes and those results would be reported in a subsequent paper.

Consider 4×4 and 10×10 2-D grids consisting of single radio (802.11b) nodes, all of which are assumed to be saturated, *i.e.*, they always have a packet to send. Each node transmits at a fixed transmit power and with equal probability to any of its grid neighbors. The *grid separation distance*, d , defined as the physical distance between any two communicating neighbors, is chosen suitably relative to the transmission range R_t . Given

a target probability of bit error rate (BER) and a corresponding SINR threshold, S_0 , which satisfies the required BER, the transmission range is given by:

$$R_t = d_{ref} \left(\frac{\bar{P}_{rx}}{S_0 P_N} \right)^{\frac{1}{\gamma}}$$

where:

- γ is the path loss exponent, typically between 2 and 4,
- P_N is the background noise power, and
- \bar{P}_{rx} is the power received at a reference point in the far field region at a distance d_{ref} from the transmitting antenna. Denoting the transmit power by P_{tx} and the wavelength by λ , the parameter \bar{P}_{rx} for $d_{ref} = 1\text{m}$ is given by:

$$\bar{P}_{rx} = \frac{P_{tx} \lambda^2}{16\pi^2}$$

For $S_0 = 11\text{ dB}$, $P_{tx} = 1\text{ mW}$ (0 dBm), $\gamma = 3$ and $P_N = -100\text{ dBm}$, the transmission range can be easily computed to be $R_t = 42.8\text{m}$. For this experiment, we have set the grid separation distance equal to half the transmission range; *i.e.*, $d = R_t/2 \approx 21\text{m}$.

The interference range, R_i , defined as the maximum distance at which the receiver corresponding to a reference transmission will be interfered with by another source (*i.e.*, the received SINR at the reference receiver drops below the threshold S_0), is given by:

$$R_i = d \left(\frac{1}{\frac{1}{S_0} - \left(\frac{d}{d_{ref}} \right)^\gamma \left(\frac{P_N}{\bar{P}_{rx}} \right)} \right)^{\frac{1}{\gamma}}$$

For $d = 21\text{m}$ and other parameters as defined above, the interference range can be computed to be $R_i \approx 52\text{m}$.

Figure 5 shows the aggregate 1-hop throughput as a function of the carrier sense range, for the 4×4 and 10×10 grids. The simulations were run with the following parameters:

- RTS/CTS mechanism disabled,
- all packets are of length 1024 bytes, and
- the sending rate is 122 packets/sec, or equivalently, 1 Mbps.

The key observation from the figure is that spatial reuse becomes possible only with higher network sizes. In the 4×4 grid, the aggregate 1-hop throughput is less than the link layer rate, implying that only one transmission occurs successfully at any time. A second observation is that, the 1-hop throughput approaches its maximum for the 10×10 grid when the carrier sense range approaches the interference range ($\approx 52\text{m}$). This is justified since a CS range, when properly tuned to the interference range, will block most potentially interfering simultaneous transmissions, thereby maximizing the aggregate 1-hop throughput.

B. One-hop Throughput Scaling as a function of Carrier Sensing Threshold

In this experiment, we investigate how the aggregate 1-hop throughput, as a function of carrier sensing range, scales with the number of orthogonal channels for a 10×10 grid. As

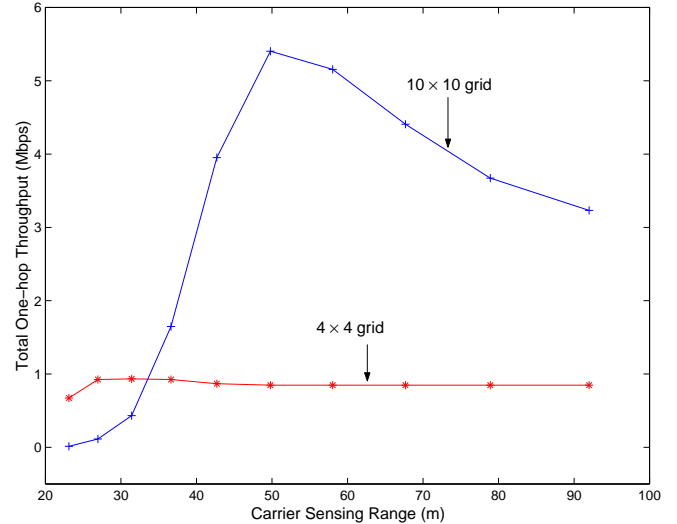


Fig. 5. Illustrating how the 1-hop total throughput (in Mbps), as a function of the carrier sensing range, scales with network size. Each node is assumed to have one 802.11b radio and only one channel is available.

in Section V-A, we concentrate on the single radio (.11b) mesh. Simulations for multi-radio meshes are currently being conducted and will be reported in a subsequent paper.

Figure 6 shows our channel assignments when 2 and 3 orthogonal channel assignments are available. Note that the assignment schemes are chosen to ensure *maximum channel diversity*, *i.e.*, each channel is used equally (on 50 radios, or equivalently, 25 links), but otherwise non-optimized. All channel assignments for this experiment were fixed; *i.e.*, dynamic channel switching was not allowed. Therefore, each node can communicate only with its neighbor with which it shares a common channel. For example, in Figure 6, node R1 can communicate with its neighbor R11 but not with R2 since they do not share a common channel.

All other simulation parameters are identical to those discussed in Section V-A.

Figure 7 shows how the aggregate 1-hop throughput, as a function of the carrier sense range, scales with the number of orthogonal channels in a 10×10 802.11b mesh grid. The results clearly show the benefits of using multiple orthogonal channels - the maximum aggregate 1-hop throughput scales (nearly) proportionally to the number of channels. This is explained by the fact that, when the CS range is equal to the interference range (at which point the throughput is maximized), the size of the collision domain is effectively halved (given the symmetry in the channel assignment scheme) when 2 channels are available, compared to when 1 channel is available. Restated, for any given link e , if the number of its neighboring potentially interfering links which are blocked by the carrier sense mechanism when one channel is available is x , the number of interfering links is approximately $x/2$ ($x/3$) when 2 (3) channels are available, thereby leading to a doubling (tripling) of the 1-hop throughput.

Moreover, when the carrier sensing range is small, the hidden terminal problem appears to be significantly alleviated with increasing number of channels, as indicated by the jump




Fig. 6. Channel assignments for simulation results discussed in Section V-A. Each node is assumed to have one 802.11b radio. (a) 2 orthogonal channels are available, (b) 3 orthogonal channels area available. Note the *symmetry* and *diversity* of the channel assignment scheme.

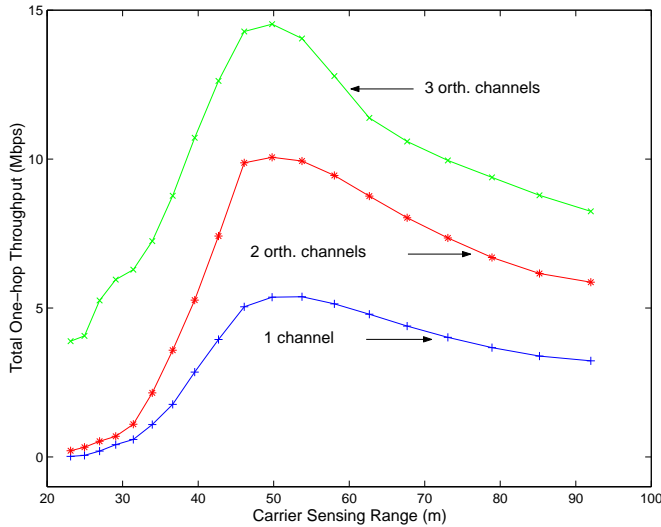


Fig. 7. Illustrating how the 1-hop total throughput (in Mbps), as a function of the carrier sensing range, scales with the number of orthogonal channels in a 10×10 802.11b mesh grid.

in throughput for 3 channels compared to 1 and 2 channels. This is intuitively justified since multiple orthogonal channels effectively shrinks the collision domain size. Although the effectiveness of the carrier sensing mechanism in avoiding collisions is reduced if the CS range is small relative to the interference range, using multiple orthogonal channels largely alleviates the collision problem, thereby leading to significantly enhanced throughput when the CS range is small but multiple channels are available.

Another observation from the figure is the presence of an optimal carrier sensing range for any number of channels, as discussed in Section V-A.

C. End-to-end Throughput with Multiple Radios and Channels: Single flow case

In this section, we investigate the various system parameters that impact the end-to-end throughput in 802.11a wireless mesh networks. We focus on the single flow case, which implies that the end-to-end throughput is dictated essentially by the extent of channel reuse (or lack thereof) in the end-to-end path. Restated, mutual interference along a path is only due to hops sharing the same channel. We are currently conducting similar simulations for multiple active flows and these results would be published in a subsequent paper. The parameters which we have used for our simulations are listed below:

- Number of nodes = 16, arranged as a regular (*i.e.*, the physical distance between any two neighbors is the same) 2-D 4×4 grid.
- Grid separation distance $d = 100\text{m}$.
- All mesh nodes are routers.
- Each node is equipped with either one or two 802.11a radios.
- Number of orthogonal data channels = 4.
- The channel assignments are as shown in Figure 8. As can be seen from the figure, some (source, destination) paths now have all their links on distinct channels (*e.g.*, $13 \rightarrow 14 \rightarrow 15 \rightarrow 16 \rightarrow 12$), whereas other paths may have all links on the same channel (*e.g.*, $3 \rightarrow 2 \rightarrow 6 \rightarrow 5 \rightarrow 9$).
- Transmission range = 150m.
- A *flow* refers to an IP connection between a source-destination pair.
- Packet size = 1500 bytes, sending rate= 1000 packets/second.
- Transmission power = 1 mW.

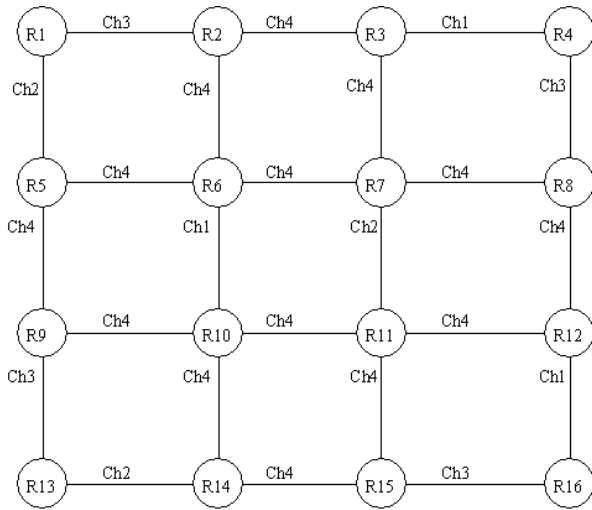


Fig. 8. Channel assignments for simulation results discussed in Section V-C. Each node is assumed to have 2 802.11a radios and 4 orthogonal channels are available.

- Carrier sense threshold = -95 dB, which is equivalent to a carrier sense range of 261m.
- RTS/CTS mechanism is disabled.
- Routing is done by a static routing table built in each router to control the active flows.

We evaluate the impact of channel assignment on the throughput, as a function of the path length (in number of hops). Table I lists the throughputs obtained for networks with (a) single radio per node and one channel and (b) 2 radios per node and 4 orthogonal channels. The throughput figures reported represent the average over 10 randomly chosen (source, destination) paths, for each path length. The link data rate is set to 12 Mbps for this experiment.

TABLE I

Throughput enhancement in multi-radio, multi-channel 802.11a based wireless mesh networks, as a function of path length. The notation H in the first column of the table represents the path length (hop count). Columns 2 and 3 show the observed throughputs (in packets/sec) for the single radio, single channel case and the 2-radios, 4-channel case respectively. Column 4 shows the percent improvement when 2 radios and 4 channels are available, compared to the single radio, single channel case.

H	S (pps): 1R1C	S (pps): 2R4C	Improvement
4	163	282	73 %
5	120	251	109 %
6	93	245	163 %

It is evident from Table I that the enhancement in throughput when multiple radios and channels are available is an increasing function of the path length. This is a consequence of improved channel reuse which is possible with multiple radios and channels.

We note that the maximum achievable throughput for a single flow is only one possible figure of merit for evaluating network performance. The channel assignment algorithm we have used is not optimized for this metric; nevertheless, the gains from the use of multiple radios are clear.

VI. CONCLUSION

In this work, we have highlighted the potential of multi-radio wireless mesh networks, along with the primary technical challenges that must be addressed for widespread deployment of such networks. Specifically, the additional degrees of freedom afforded by having multiple radios per mesh node in scaling both the aggregate 1-hop and end-to-end throughput was highlighted. The approaches required are necessarily ‘cross-layer’; for example, suitable channel-aware metrics are critical to solving the channel assignment problem. We anticipate such multi-radio mesh networks to be a focus of continuing research and evaluation driven by developments in the 802.11s Task Group that is seeking to currently define Layer2 attributes in support of these goals.

REFERENCES

- [1] J. Zhu and S. Roy, “802.11 mesh networks with two radio access points,” *Proc. International Communications Conference*, 2005.
- [2] P. Gupta and P. R. Kumar, “The capacity of wireless networks,” *IEEE Trans. Info. Theory*, vol. 46, pp. 388–404, March 2000.
- [3] J. Li, C. Blake, D.S.J. De Couto, H.I Lee, and R. Morris, “Capacity of ad hoc wireless networks,” *Proc. MOBICOM*, 2001.
- [4] D. S. J. De Couto, D. Aguayo, B. A. Chambers and R. Morris, “Performance of Multihop Wireless Networks: Shortest path is not Enough,” *Proc. HotNets I*, Oct. 2002.
- [5] Ashish Raniwala, Kartik Gopalan, and Tzi cker Chiueh, “Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks,” *SIGMOBILE Mobile Computing and Communications Review*, vol. 8, no. 2, pp. 50–65, 2004.
- [6] P. Kyasanur and N. Vaidya, “Routing and interface assignment in multi-channel multi-interface wireless networks,” *Proc. of IEEE WCNC*, 2005.
- [7] Richard Draves, Jitendra Padhye, and Brian Zill, “Routing in multi-radio, multi-hop wireless mesh networks,” in *MobiCom '04: Proceedings of the 10th Annual International Conference on Mobile Computing and Networking*, New York, NY, USA, 2004, pp. 114–128, ACM Press.
- [8] P. Gupta, R. Gray, and P. R. Kumar, “An experimental scaling law for wireless ad-hoc networks,” *Technical Report, University of Illinois at Urbana-Champaign*, May 2001.
- [9] Ashish Raniwala and Tzi cker Chiueh, “Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network,” *Proc. of IEEE INFOCOM*, 2005.
- [10] X. Guo, S. Roy, and W. S. Conner, “Spatial reuse in wireless ad-hoc networks,” *Proc. Vehicular Technology Conference*, 2003.
- [11] J. Zhu, X. Guo, L. Yang, W. Steven Conner, S. Roy, and M. Hazra, “Adapting physical carrier sensing to maximize spatial reuse in 802.11 mesh networks,” *Wiley J. Wireless Communications and Mobile Computing, Spl. Issue on Emerging WLAN Applications and Technologies*, vol. 4, no. 8, pp. 933–946, 2004.
- [12] C. R. Lin and M. Gerla, “Adaptive clustering for mobile wireless networks,” *IEEE Jour. Selected Areas in Communications*, pp. 1267–1275, Sept. 1997.
- [13] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, “Impact of interference on multi-hop wireless network performance,” *Proc. of ACM Mobicom*, Sept. 2003.
- [14] Jungmin So and Nitin Vaidya, “Multi-channel mac for ad hoc networks: Handling multi-channel hidden terminals using a single transceiver,” *Proc. of MOBIHOC*, May 2004.
- [15] P. C. Ng, S. C. Liew, and L. B. Jiang, “Achieving scalable performance in large-scale ieee 802.11 wireless networks,” *IEEE Wireless Comm. and Networking Conf.*, Mar. 2005.
- [16] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, “Performance anomaly of 802.11b,” *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, vol. 2, no. 30, pp. 836–843, 2004.
- [17] P. Belanger, “Wireless meshmesh infrastructure - can it work for large scale deployments,” <http://www.belainetworks.com>.
- [18] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, “High throughput path metric for multihop wireless routing,” *Proc. of ACM Mobicom*, Sept. 2003.
- [19] R. Draves, J. Padhye and B.D. Zill, “Comparison of routing metrics for static multi-hop wireless networks,” *Proc. of ACM Sigcomm*, Oct. 2004.

- [20] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," *Microsoft Technical Report*, vol. MSR-TR-2003-44, July 2003.
- [21] P. Bahl, A. Adya, J. Padhye, and A. Wolman, "Reconsidering wireless systems with multiple radios," *ACM Sigcomm Comp. Commun. Review*, pp. 39–46, Oct. 2004.
- [22] P. Kyasanur and N. Vaidya, "Capacity of multichannel wireless networks: Impact of number of channels and interfaces," *Tech. Report, University of Illinois at Urbana Champaign*, 2005.
- [23] J. Zhu, S. Roy, X. Guo, and W. Steven Conner, "Maximizing aggregate throughput in 802.11 mesh networks with physical carrier sensing and two-radio multichannel clustering," *Proc. NSF-RPI Workshop on Pervasive Computing and Networking*, 2004.
- [24] B. Raman and K. Chebrolu, "Revisiting MAC Design for 802.11-based Mesh Networks," *HotNets III*, San Diego, CA, Nov. 2004.
- [25] L. Iannone, R. Khalili, K. Salamatian and S. Fdida, "Cross Layer Routing in Wireless Mesh Networks," *1st Int. Symposium on Wireless Commn. Syst.*, Sep. 2004, pp. 319-323.