

Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks

Krishna N. Ramachandran, Elizabeth M. Belding, Kevin C. Almeroth, Milind M. Buddhikot[†]
University of California at Santa Barbara[†] Lucent Bell Labs, Holmdel
{krishna, ebelding, almeroth}@cs.ucsb.edu mbuddhikot@lucent.com

Abstract—The capacity problem in wireless mesh networks can be alleviated by equipping the mesh routers with multiple radios tuned to non-overlapping channels. However, channel assignment presents a challenge because co-located wireless networks are likely to be tuned to the same channels. The resulting increase in interference can adversely affect performance. This paper presents an interference-aware channel assignment algorithm and protocol for multi-radio wireless mesh networks that address this interference problem. The proposed solution intelligently assigns channels to radios to minimize interference within the mesh network and between the mesh network and co-located wireless networks. It utilizes a novel interference estimation technique implemented at each mesh router. An extension to the conflict graph model, the multi-radio conflict graph, is used to model the interference between the routers. We demonstrate our solution’s practicality through the evaluation of a prototype implementation in a IEEE 802.11 testbed. We also report on an extensive evaluation via simulations. In a sample multi-radio scenario, our solution yields performance gains in excess of 40% compared to a static assignment of channels.

I. INTRODUCTION

Typical deployments of static multi-hop wireless networks, called *wireless mesh networks*, utilize routers equipped with only one IEEE 802.11 radio. IEEE 802.11 radios are typically *single-channel* radios. As a result, single-radio mesh networks can suffer from serious capacity degradation due to the half-duplex nature of the wireless medium [10].

Fortunately, the IEEE 802.11 PHY specification permits the simultaneous operation of multiple *non-overlapping* channels. For example, three non-overlapping channels in the 2.4GHz band can be simultaneously used. The IEEE 802.11a specification allows up to twelve non-overlapping channels in the 5.0 GHz band. By deploying multi-radio routers in wireless mesh networks and assigning the radios to non-overlapping channels, the routers can communicate simultaneously with minimal interference in spite of being in direct interference range of each other. Therefore, the capacity of wireless mesh networks can be increased.

In equipping routers with multiple radios, a naïve strategy would be to equip each router with the number of radios equal to the number of orthogonal channels. However, this strategy is economically prohibitive due to the significant number of non-overlapping channels. Furthermore, small form-factor embedded systems used for manufacturing routers support only a limited number of radios. Consequently, using all non-overlapping channels on a mesh router is still not a viable option.

The assignment of channels to a mesh router then becomes a problem of choosing which channels to assign to which of its radios. A simple technique is to use *static channel assignment*. However, with the explosive growth in “WiFi” deployments that operate in the same (unlicensed) spectrum as wireless mesh networks, any static assignment will likely result in the operation of the mesh on channels that are also used by co-located WiFi deployments. The resulting increase in interference can degrade the performance of the mesh network.

This paper addresses the channel assignment problem and specifically investigates the *dynamic* assignment of channels in a wireless mesh network. We present a centralized, interference-aware channel assignment algorithm and a corresponding channel assignment protocol aimed at improving the capacity of wireless mesh networks by making use of all available non-overlapping channels. The algorithm intelligently selects channels for the mesh radios in order to minimize interference within the mesh network and between the mesh network and co-located wireless networks. Each mesh router utilizes a novel interference estimation technique to measure the level of interference in its neighborhood because of co-located wireless networks. The algorithm utilizes an extension to the conflict graph model [14], the Multi-radio Conflict Graph (MCG), to model interference between the multi-radio routers in the mesh. The MCG is used in conjunction with the interference estimates to assign channels to the radios.

One potential pitfall of dynamic channel assignment is that it can result in a change in the network topology. Topology changes can lead to sub-optimal routing and even network partitioning in case of node failures. The proposed solution, therefore, ensures that channel assignment does not alter the network topology by mandating that one radio on each mesh router operate on a *default channel*. A second potential pitfall is that channel assignment can result in disruption of flows when the mesh radios are reconfigured to different frequencies. To prevent flow disruption, *link redirection* is implemented at each mesh router. This technique redirects flows over the default channel until the channel assignment succeeds.

We evaluate our proposed solution through simulations in Qualnet. We utilize the Optimized Link State Routing (OLSR) protocol [8] and the Weighted Cumulative Expected Transmission Time (WCETT) metric [9] for route selection. We demonstrate the practicality of our proposed solution via the evaluation of a prototype implementation in a multi-radio IEEE 802.11b testbed.

A. Research Contributions

To the best of our knowledge, ours is the first solution to address the problem of dynamic channel assignment in wireless mesh networks in the presence of interference from co-located wireless networks. A key goal in the design of the proposed solution has been to make the solution *amenable to easy implementation* using *currently available* radios. This differentiates our work from several proposed solutions (surveyed in Section IX) which require either specialized as yet unavailable radios or knowledge about the network such as anticipated traffic patterns and the specific paths to be traversed by network flows.

Specifically, the contributions of this paper are as follows:

- A dynamic, interference-aware channel assignment algorithm that minimizes interference between the mesh network and co-located wireless networks.
- A multi-radio conflict graph, an extension to the well-known conflict graph model, to model the interference relationship between multi-radio routers in a wireless mesh network.
- A novel interference estimation scheme that routers use to estimate the interference level in their neighborhoods.
- A link redirection protocol that prevents the disruption of flows during channel assignment.
- A comprehensive performance study that shows significant throughput improvements in the presence of varying interference levels, which are validated through empirical measurements on a prototype implementation.

B. Paper Outline

The remainder of the paper is organized as follows: Section II discusses the effect of channel assignment on network topology. In Section III, we formulate the channel assignment problem. Section IV describes our interference estimation technique and the multi-radio conflict graph model. In Section V, we present our centralized channel assignment algorithm. We discuss the challenges we addressed during the development of our prototype implementation in Section VI. Section VII presents results from our simulation-based evaluation, while results from our prototype evaluation are presented in Section VIII. In section IX, we summarize related work, and Section X concludes the paper.

II. CHANNEL ASSIGNMENT AND NETWORK TOPOLOGY

In a multi-radio mesh network, channel assignment to radios can alter the network topology. Consider the example four node topology in Figure 1(a). Here, node C is equipped with three radios and the other nodes (A, B, and D) have one radio each. Each link in the figure is labeled with its channel number. Figure 1(a) illustrates the topology when all radios are tuned to channel one. Figure 1(b) illustrates the change in network topology after channel assignment.

Alterations in the network topology have three main drawbacks. First, subsequent node failures have a higher probability of causing network partitions. Consequently, portions of the mesh may become unreachable, resulting in the disruption




Fig. 1. Network topology with varying channel assignments.

of flows. This is clearly seen in Figure 1(c). When node C fails, the four node network is partitioned into three different clusters. Reconnection of the network would require complex synchronization schemes to be implemented at the mesh routers [25].

Second, topology alterations can result in *sub-optimal routes* between node pairs with respect to some metric, such as throughput, delay, or reliability. To illustrate how this can occur, consider again Figure 1(a). Node A can communicate with node B on a one hop path. After channel assignment, A can communicate with B only over a two hop path via C, as shown in Figure 1(b). Selection of a path with a higher hop-count is not preferred for three reasons: (1) longer, frequency diversified paths often yield worse performance than shorter paths; (2) the interference “foot-print” of a flow on a higher hop-count path is naturally greater; and (3) a longer path is more prone to failure. Note, however, that we do not claim that all longer paths are likely to perform poorly compared to shorter paths because the performance of each path alternative is likely to vary with factors such as traffic pattern, node placement, radio characteristics, and terrain. Nevertheless, we stress that it is challenging to accurately predict, in practice, which channel assignment alternative and resulting network topology configuration will yield optimum performance.

The third drawback of altering a network’s topology is that it affects existing flows. For example, let us assume that link *CD* in Figure 1(b) is assigned a new channel. The process of channel assignment must be accurately coordinated; otherwise, cases may arise where one radio on the link switches to the new frequency but the second radio does not because a control message is either lost or delayed. Consequently, any flows from *D* to the rest of the network that existed at the time of the channel assignment are disrupted during the switch. Overcoming such cases is challenging in practice because configuration of the radios requires time-synchronized coordination between the mesh routers during channel assignment.

Because of these drawbacks associated with network topology changes, we advocate that topology alterations should be avoided. We mandate this by requiring that all routers in the mesh network designate one of their radios to be a *default radio interface*. This default radio is of the same physical layer technology, either 802.11a, 802.11b, or 802.11g, and is tuned to a common channel throughout the mesh. The *default channel* carries both control and data traffic.

This strategy has several advantages. First, it prevents changes in the topology of the network because routers will discover otherwise disconnected neighbors by communicating over the default radio interface. Second, overcoming node

failure is simplified because a router will be able to choose alternate paths to route around a failed node. Third, the routing protocol will now have the option of selecting a path that is not frequency diversified if it has better performance characteristics than a frequency diversified alternative. As a final advantage, any disruption of flows during channel assignment can be avoided by *redirecting* flows over the default radio until the assignment completes. The redirection technique is further elaborated in Section VI. We consider the reassignment of the default channel in Section VI-D.

III. PROBLEM FORMULATION

The channel assignment algorithm we propose in this paper is designed for wireless mesh networks. Routers in such networks are stationary. However, user devices, such as laptops and PDAs, can be mobile. Such devices associate with routers that also function as access points.

Figure 2 illustrates our model of a multi-radio mesh network. In our model, the mesh routers are assumed to be equipped with multiple IEEE 802.11 radios, such as 802.11a, 802.11b, or 802.11g. The routers need not all be equipped with the same number of radios nor do they need all three types of radios. Depending on the number of radios at each mesh router, we classify the routers into two categories: (1) Multi-Radio mesh routers (MRs); and (2) Single-Radio mesh routers (SRs). We mandate that each MR and SR in the network be equipped with one radio, called the *default radio*, which is of the same physical layer type, e.g. 802.11b, and tuned to the same channel as motivated in Section II.

At least one router in the mesh is designated as a *gateway*. The gateway provides connectivity to an external network. In order to simplify the explanation of the channel assignment solution, we assume the presence of only one gateway. Access Points (APs) provide connectivity to user devices and are co-located with mesh routers. A majority of the traffic within the mesh is either from the user devices to the gateway or vice-versa. This traffic pattern is typical in wireless mesh deployments. Because the traffic pattern is skewed to-and-from the gateway, the paths taken by the resulting flows are likely to form a tree structure in which the gateway is the “root” and the user devices are the “leaves”. Traffic flows will likely aggregate at routers close to the gateway. Therefore, in order to improve overall network capacity, it is preferable to place MRs close to the gateway and in regions of the mesh that are likely to experience heavy utilization. It is important that the placement occur after careful *network planning* in order to optimize network performance, reduce equipment costs, and address logistical constraints.

The dotted lines in the figure illustrate links between MRs that are tuned to non-overlapping channels. In our example, five such channels are used. A sixth channel, indicated by solid lines, is the default channel. The Channel Assignment Server (CAS), which is co-located with the gateway in the figure, performs channel assignment to radios.

In assigning channels, the CAS should satisfy the following goals:

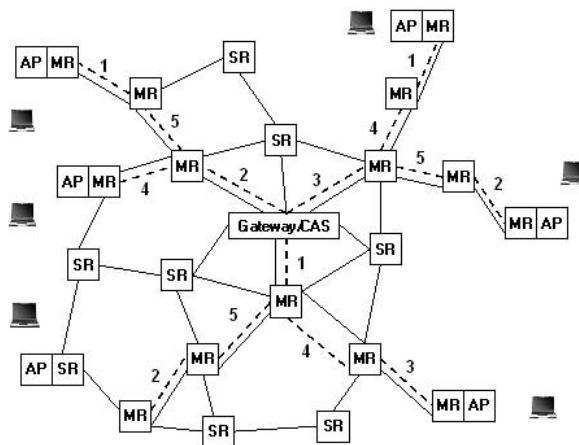


Fig. 2. Multi-Radio wireless mesh architecture.

- *Minimize interference between routers in the mesh:* In satisfying this goal, three sub-goals need to be achieved. First, the CAS should satisfy the constraint that for a link to exist between two routers, the two end-point radios on them must be assigned a common channel. Second, links in direct communication range of each other should be tuned to non-overlapping channels. Third, because of the tree shaped traffic pattern expected in wireless mesh networks, channel assignment priority should be given to links starting from the gateway and then to links fanning outwards towards the edge of the network.
- *Minimize interference between the mesh network and wireless networks co-located with the mesh:* In satisfying this goal, the CAS should periodically determine the amount of interference in the mesh due to co-located wireless networks. The interference level is estimated by individual mesh routers. The CAS should then re-assign channels such that the radios operate on channels that experience the least interference from the external radios.

Given these goals for the channel assignment algorithm, next we present details on interference estimation and describe the interference modeling technique.

IV. INTERFERENCE ESTIMATION AND MODELING

This section presents an overview of the interference estimation procedure. Implementation details are left to Section VI. This section also introduces the *Multi-radio Conflict Graph* (MCG) model.

A. Interference Estimation

The goal of interference estimation is to periodically measure the interference level in each mesh router’s environment. Accurate measurement, however, is challenging and requires that expensive hardware be used [1].

Instead, as an approximation, we rely on the number of *interfering radios* on each channel supported by each router as an estimation of interference. An interfering radio is defined as a simultaneously operating radio that is *visible* to a router

but *external* to the mesh. A visible radio is one whose packet(s) pass Frame Check Sequence (FCS) checks and are therefore correctly received. We assume that the CAS informs the router of radios *internal* to the mesh. The information could consist of an IP address range or an exhaustive list of all radio MAC addresses in the mesh.

One caveat to the above estimation procedure is that *carrier-sensing* radios, i.e., those radios that are within an estimating router's carrier sensing range but outside its reception range, will *not* be accounted for in the estimation. This is because packets transmitted by such radios will fail FCS checks performed by the router. However, carrier-sensing radios may still interfere with the router. Our interference estimation technique does not consider such radios for two reasons. First, recent studies [12], [24] suggest that current IEEE 802.11 MAC implementations are overly conservative in their carrier sense mechanism and often overestimate the adverse impact of interfering radios. Therefore, even in the presence of multiple carrier-sensing radios, the performance degradation due to carrier-sensing neighbors may not be as severe as previously understood. Second, even if we were to incorporate carrier-sensing radios in our interference estimation solution, it is impossible to determine the presence of such radios using commodity hardware because of the inability of current firmware implementations to identify them¹. A recent proposal [22] aims to overcome the firmware limitations by using specialized hardware². Such hardware are likely to be available in the future. Another proposal aims to discover carrier-sensing neighbors using pairwise broadcast probing [17]. However, it has the drawback that the probing procedure can take a long time to complete. Therefore, utilizing it in our solution is still not feasible. We are currently investigating strategies to address its drawback. In the meantime, we assume a solution exists, and we leverage this assumption in our implementation.

Measurement of only the number of interfering radios, however, is not sufficient because it does not indicate the amount of traffic generated by the interfering radios. For instance, two channels could have the same number of interfering radios but one channel may be heavily utilized by its interfering radios compared to the other. Therefore, in addition, each mesh router also estimates the channel bandwidth utilized by the interfering radios.

The interference estimation procedure is as follows: a mesh router configures one radio of each supported physical layer

¹Wireless devices, such as ones using the Prism 2/2.5 chipset, sometimes allow the capture of packets transmitted by carrier-sensing radios that fail the FCS check. This mechanism at first might suggest a technique to identify the carrier-sensing radios. However, the utility of this capture mechanism is limited because the information contained in the garbled packets is by nature faulty.

²An approach that avoids specialized hardware assumes that the carrier-sensing range is k times the reception range [19]. We note that the relationship is non-deterministic and less likely to be effective in practice because the carrier-sensing range is dependent on a myriad number of factors such as transmission power, receiver sensitivity, environmental conditions, and the presence of obstacles.

type to capture packets³ on each supported channel for a small duration. The router uses the captured packets to measure the number of interfering radios and per second channel utilization. The number of interfering radios is simply the number of unique MACs external to the mesh. The utilization on each channel due to the interfering radios is computed from the captured data frames by taking into account the packet sizes and the rates at which the packets were sent [13]. The overhead of the MAC layer is accounted for in our utilization calculation. We set the duration of the packet capture to three seconds in our implementation. The three second duration is large enough to allow for the averaging of the variations in per second measurements and is small enough to enable the interference estimation to complete quickly.

Each mesh router then derives two separate channel rankings. The first ranking is according to increasing number of interfering radios. The second ranking is according to increasing channel utilization. The mesh router then merges the rankings by taking the average of the individual ranks. The resulting ranking is sent to the CAS.

B. Interference Modeling

Conflict graphs are used extensively to model interference in cellular radio networks [14]. A conflict graph for a mesh network is defined as follows: consider a graph, G , with nodes corresponding to routers in the mesh and edges between the nodes corresponding to the wireless links. A conflict graph, F , has vertices corresponding to the links in G and has an edge between two vertices in F if and only if the links in G denoted by the two vertices in F interfere with each other. As an example of a conflict graph, Figure 3(a) shows the topology of a network with four nodes. Each node in the figure is labeled with its node name and its number of radios. Figure 3(b) shows the conflict graph.

At a first glance, the problem of assigning channels to links in a mesh network appears to be a problem of vertex coloring the conflict graph. However, vertex coloring fails to assign channels correctly because it does not account for the constraint that the number of channels (colors) assignable to a router must be equal to its number of radios. As an example of why this is the case, let us assume that the four vertices in the conflict graph shown in Figure 3(b) are each assigned one of three different channels using a vertex coloring algorithm. This means that the two radios represented by each vertex in the conflict graph operate on the frequency assigned to that vertex. This implies that node C in the illustrated network operates on three different channels, which is impossible because it is equipped with only two radios.

The conflict graph does not correctly model routers equipped with multiple radios. Therefore, we extend the conflict graph to model multi-radio routers. In the extended

³Packet capture mode as implemented on currently available IEEE 802.11 radios cannot capture packets from radios, such as cordless phones or Bluetooth devices, that use other physical layer technologies. We note, however, that the interference foot-print of such devices is likely to be small. Software-defined radios are likely to address this limitation in the future.




Fig. 3. (a) A simple network topology, G . (b) Corresponding conflict graph, F' . (c) Corresponding multi-radio conflict graph, F' .

model, called the Multi-radio Conflict Graph (MCG), we represent edges between the mesh radios as vertices instead of representing edges between the mesh routers as in the original conflict graph. To create the MCG, F' , we first represent each *radio* in the mesh network as a vertex in G' instead of representing routers by vertices as in G . Therefore, in the above example, node C is represented by two vertices in G' corresponding to its two radios instead of just one vertex in G . The edges in G' are between the mesh radios instead of the mesh routers as in G . We then represent each edge in G' using a vertex in F' . The edges between the vertices in F' are created in the same way the original conflict graph is created, i.e., two vertices in F' have an edge between each other if the edges in G' represented by the two vertices in F' interfere with each other. As an example, Figure 3(c) shows the multi-radio conflict graph of the network shown in Figure 3(a). In the figure, each vertex is labeled using the radios that make up the vertex. For example, the vertex $(A - 1 : C - 2)$ represents the link between the first radio on router A and the second radio on router C .

When using a vertex coloring algorithm to color the MCG, we impose an important constraint: on coloring any MCG vertex, all uncolored vertices in the conflict graph that contain any radio from the just-colored vertex be removed. For example, after assigning a color to vertex $(A - 1 : C - 2)$ in Figure 3(c), all vertices containing either $A - 1$ or $C - 2$ should be removed from the conflict graph. This is required to ensure that only one channel is assigned to each radio in the mesh network.

V. CHANNEL ASSIGNMENT ALGORITHM

A. Overview

The channel assignment problem for mesh networks is similar to the *list coloring* problem, which is defined as follows: given a graph, $G = (V, E)$, and for every v in V , a list $L(v)$ of colors, is it possible to construct a valid vertex coloring of G such that every vertex v receives a color from the list $L(v)$? The list coloring problem is NP-complete [21]. Therefore, we rely on an approximate algorithm for channel assignment. Our algorithm, called the Breadth First Search Channel Assignment (*BFS-CA*) algorithm, uses a breadth first search to assign channels to the mesh radios. The search begins with links emanating from the gateway node. The rationale behind the use of breadth first search is intuitive: by using breadth first search, we satisfy our goal described in Section III

of giving channel assignment priority to links starting from the gateway and then in decreasing levels of priority to links fanning outward towards the edge of the network.

Before using the *BFS-CA* algorithm, the channel assignment server (CAS) obtains the interference estimates from the mesh routers. It then chooses a channel for the default radios. The default channel is chosen such that its use in the mesh network minimizes interference between the mesh network and co-located wireless networks. The CAS then creates the MCG for the non-default radios in the mesh. The MCG is created using the neighbor information sent by each mesh router to the CAS. After constructing the MCG, the CAS uses the *BFS-CA* algorithm to select channels for the non-default radios. Once the channels are selected for the mesh radios, the CAS instructs the routers to configure their radios to the newly selected channels. To simplify the explanation of the channel selection procedure in this section, let us assume for now that the mesh radios are reconfigured at the same time. We address this assumption in Section VI-D, where we provide details on the specific protocol used to re-assign channels.

The default channel selection procedure is presented next followed by a detailed description of the *BFS-CA* algorithm. The CAS periodically invokes the channel selection procedure summarized above to cope with the varying nature of interference in the mesh. This section ends with a discussion of the period of invocation and its implications.

B. Default Channel Selection

The CAS chooses the default channel using the rank of a channel, c , for the entire mesh, R_c . R_c is computed as follows:

$$R_c = \frac{\sum_{i=1}^n Rank_c^i}{n}$$

where n is the number of routers in the mesh and $Rank_c^i$ is the rank of channel c at router i . The default channel is then chosen as the channel with the least R_c value. The intuition behind this metric is to use the least interfered channel as the default channel in the mesh. Using such a channel satisfies our goal of minimizing interference between the mesh and co-located wireless networks.

C. Non-Default Channel Selection

In this phase, the CAS uses the neighbor information collected from all routers to construct the MCG. Neighbor information sent by a router contains the identity of its neighbors, delay to each neighbor, and interference estimates for all channels supported by the router's radios. Section VI-C details the calculation of link delay performed by mesh routers. The CAS associates with each vertex in the MCG its corresponding link delay value. The CAS also associates with each vertex a channel ranking derived by taking the average of the individual channel rankings of the two radios that make up the vertex. The average is important because the assignment of a channel to a vertex in the MCG should take into account the preferences of both end-point radios that make up the vertex.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.