

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:

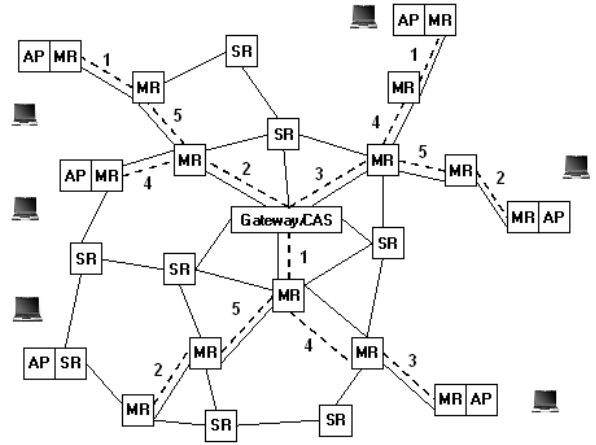


Fig. 2. Multi-Radio wireless mesh architecture.

- *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. Sanzgiri et al. propose to use specialized hardware to overcome the firmware limitations [22]. Such hardware are likely to be available in the future and can be leveraged when available.

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 type to capture packets² 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

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.

²Packet 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.

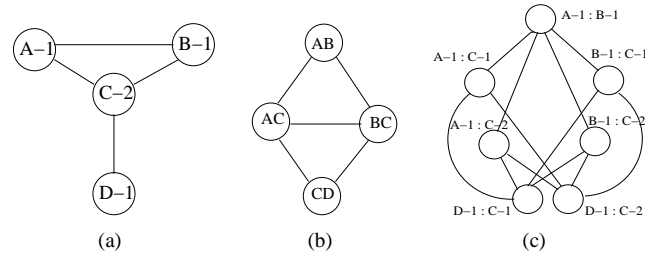


Fig. 3. (a) A simple network topology, G . (b) Corresponding conflict graph, F . (c) Corresponding multi-radio conflict graph, F' .

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 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 vertices 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. We use a two hop interference model to create the MCG. In this model, two mesh links are interfering if they either have a common router or are separated by a hop as determined from the neighbor information sent by each mesh router to the CAS. Padhye et

al. propose an empirical technique to detect interfering mesh links [17]. This technique can be more accurate because of its empirical nature, however, takes a long time (several hours) to complete. We are currently investigating approaches to speed up this technique. In the meantime, we leverage the two hop model in our work.

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.

For all vertices in the MCG, the CAS then computes their distances from the gateway. The distance of an MCG vertex is the average of the distances from the gateway of the two radios that make up the vertex. The distance of a radio is obtained from beacons initiated by the gateway. A beacon is a *gateway advertisement* broadcasted hop-by-hop throughout the mesh.

Algorithm 1 BFS-CA Algorithm

```
1: Let  $V = \{v|v \in \text{MCG}\}$ 
2: while notAllVerticesVisited $\{V\}$  do
3:   Let  $h = \text{smallestHopCount}(V)$ 
4:    $Q = \{v|v \in V \text{ and notVisited}(v) \text{ and hopcount}(v) == h\}$ 
5:   sort( $Q$ )
6:   while size( $Q$ ) > 0 do
7:      $v_{\text{current}} = \text{removeHead}(Q)$ 
8:     if visited( $v_{\text{current}}$ ) then
9:       continue
10:    end if
11:    visit( $v_{\text{current}}$ )
12:     $V_n = \{u|u \in \text{MCG} \text{ and edgeInMCG}(u, v_{\text{current}}) == \text{TRUE}\}$ 
13:    permanently assign highest ranked channel  $c$  from  $v_{\text{current}}$ 's channel ranking that does not conflict with  $u_i$ ,  $\{u_i \in V_n \text{ and } 0 \leq i < \text{size}(V_n)\}$ 
14:    if  $c$  does not exist then
15:      permanently assign random channel to  $v_{\text{current}}$ 
16:    end if
17:     $L = \{v|v \in \text{MCG} \text{ and } v \text{ contains either radio from } v_{\text{current}}\}$ 
18:    removeVerticesInListFromMCG( $L$ )
19:    tentatively assign  $c$  to radios in  $L$  that are not part of  $v_{\text{current}}$ 
20:    Let  $r_f$  be router with interface in  $v_{\text{current}}$  that is farthest away from gateway
21:    Let  $Tail =$  list of all active  $v$  ( $v \in \text{MCG}$ ) such that  $v$  contains an interface from  $r_f$ 
22:    sort( $T$ )
23:    addToQueue( $Q, Tail$ )
24:  end while
25:  permanently assign channels to radios that are unassigned a permanent channel.
26: end while
```

Each beacon contains a hop-count field that is incremented at each hop during its broadcast. The distance of a router from the gateway is the shortest path length of a single beacon instance received by the router over all paths. The router communicates the distance to the CAS via periodic *heartbeat* messages sent every minute in our implementation.

Algorithm: Once the average distances are computed, the CAS uses the *BFS-CA* algorithm to assign channels to the mesh radios. The algorithm is summarized in Algorithm 1.

The algorithm starts by adding all vertices from the MCG to a list, V (Line 1). It does a breadth first search of the MCG to visit all vertices and assign them channels. The search starts from vertices that correspond to links emanating from the gateway (Lines 3, 4). In line 3, the smallest hopcount vertex is determined of all vertices in the MCG. In line 4, all vertices with distance equal to the smallest hop count are added to a queue, Q . If vertices correspond to network links emanating from the gateway, their hop count is 0.5. These vertices are then sorted by increasing delay values (Line 5). This sort is performed in order to give higher priority to the better links emanating from the shortest hop count router (the gateway for the first BFS iteration).

The algorithm then visits each vertex in Q (Line 11) and *permanently* assigns them the highest ranked channel that does not conflict with the channel assignments of its neighbors (Line 13). If a non-conflicting channel is not available, a randomly chosen channel is permanently assigned to the vertex. Note, however, that the default channel is never assigned. Once a vertex is assigned a channel, all vertices that contain either radio from the just-assigned vertex are placed in a list, L (Line 17). In line 18, all vertices from L are removed from the

MCG. This step is needed to satisfy the constraint that only one channel is assigned to each radio. The radios in the list of vertices that do not belong to the just-assigned vertex are *tentatively* assigned the latter's channel (Line 19).

In lines 20-21, vertices at the next level of the breadth first search are added to Q . These vertices correspond to links that fan-out from the gateway towards the periphery. To find such links in the MCG, two steps are performed. In the first step (Line 20), the router from the just-assigned vertex that is farthest away from the gateway is chosen; the farthest router is the router with the higher hop-count of the two routers that make up the just-assigned vertex. In the second step (Line 20), all unvisited MCG vertices that contain a radio belonging to the farthest router are added to the list, $Tail$. This list is sorted (Line 22) by increasing value of the delay metric to give higher priority to better links that emanate from the farthest router. Finally, in line 23, the vertices from $Tail$ are added to Q . The above described algorithm continues until all vertices in the MCG are visited. In line 25, any radio that is not assigned a permanent channel during the search, because vertices containing it were deleted in line 18, is permanently assigned one of the channels tentatively assigned to it in line 19.

Once channel assignments are decided, the CAS notifies the mesh routers to re-assign their radios to the chosen channels. The exact protocol is described in Section VI-D.

D. Channel Re-assignment Strategy

To adapt to the changing interference characteristics, the CAS periodically re-assigns channels. The periodicity depends ultimately on how frequently interference levels in the mesh network are expected to change. If a large number of interfering devices in the vicinity of the mesh network are expected to be short-lived, the invocation rate should be increased. On the other hand, if a majority of the interfering devices are likely to be long-lived, the invocation rate can be decreased. In our implementation, we have set the rate to ten minutes. We believe this rate results in a good tradeoff between interference adaptation and mesh radio reconfiguration. Nevertheless, we expect the network operator of a mesh network to choose a rate to best suit the target deployment.

VI. IMPLEMENTATION CONSIDERATIONS

In the earlier sections, we omitted details on several implementation specific steps, such as the channel estimation procedure, link delay estimation, and the channel assignment process. This section describes the implementation aspects of these steps as they become critical when *BFS-CA* is used in practice.

A. Interference Estimation

Each mesh router performs interference estimation by capturing packets from the medium. To capture packets, we use a special operating mode supported by typical radio hardware called the *RFMon* mode. This mode allows IEEE 802.11 management frames and regular data frames in the medium to

be captured. The management frames are periodically transmitted (typically every 100 msec) by devices that implement the IEEE 802.11 specification. We do not use the typical promiscuous mode supported by the radios because that mode only sniffs data frames and does not sniff IEEE management frames. Consequently, devices that do not transmit any data packets during the capture period will not be discovered. The RFMon mode provides information about each packet, such as the rate at which it was sent and the size of the packet. It derives this information by interpreting physical layer information associated with the packet and therefore can provide the information even if layer 2 encryption is used to conceal the contents of the packets. The captured packets are then used to measure the number of interfering radios and the bandwidth consumed by those radios.

Radios, when placed in RFMon mode, cannot transmit data packets for the duration of the sniffing. This is because some commercial radios, such as those utilizing the Prism 2/2.5 chipset, cannot perform packet transmission in RFMon mode. Even if radios can transmit in RFMon mode, such as ones utilizing the Atheros chipset, because the interference estimation at each router occurs independently of its neighbors, there is no guarantee that two radios will be on the same channel in order to communicate successfully. Therefore, any flow that uses a radio temporarily in RFMon mode will be disrupted. To prevent flow disruption, we use *link redirection*, described next.

B. Link Redirection

Link redirection is achieved when a flow intended for a router's non-default radio is redirected to the router's default radio instead. Link redirection is invoked in two cases: (1) when a router's intended transmitter is incapable of delivering packets, and (2) when the intended receiver on the neighboring router is incapable of receiving packets. Redirection is possible because the default radios on all routers operate on the default channel.

Our link redirection protocol is as follows: whenever a radio has to change its state to inactive, it broadcasts an INTERFACE-INACTIVE message every second for three seconds before it changes state. The multiple broadcasts allow for any message losses. Any neighbor that receives the INTERFACE-INACTIVE message deletes the address of the soon-to-become inactive radio from its routing tables. Once the deletion occurs, link redirection is invoked until the radio becomes active again. We simply rely on the routing protocol to notify that a radio has become active.

C. Link Delay Estimation

We measure link delay using the Expected Transmission Time (*ETT*) metric [9]. *ETT* of a link is derived from the link's bandwidth and loss rate. Due to space constraints, we specify below only the values of various parameters used to calculate *ETT*. A more detailed description of the metric can be found in [9]. *ETT* of a link is given as $(etx * s/b)$ where *etx* is the expected number of transmissions necessary to send a

packet on the link; *s* is the size of the packet (set to 1024 bytes in our implementation); and *b* is the bandwidth of the link. For the *etx* metric, we set the broadcast probe size to 1024 bytes and probe rate to one second. *etx* is calculated every 10 seconds. The link data rate, *b*, is determined using packet pair probing [15]. The packet pair is of size 134 bytes and 1200 bytes respectively and is sent every minute. We choose the minimum of ten samples to estimate *b*.

D. Channel Assignment Protocol

Channel assignment is different for non-default radios and default radios. Assignment of non-default radios can happen without flow disruption by invoking link redirection. To perform the assignment, the CAS simply sends a message to the router whose radios need to be reconfigured. The CAS waits for a positive acknowledgment from the router for a small time interval (5 seconds in our implementation). If the router fails to acknowledge the message, the CAS sends the message up-to five times in our implementation. As soon as a router receives the channel assignment message, it invokes link redirection by following the protocol described in Section VI-B. It then initiates configuration of its radios. After the assignment completes and the neighboring radio on the link is confirmed to be active, redirection stops. In our implementation, we use the *ping* utility to confirm that the link is active.

In order to assign the default radio channel, we assume the availability of a reliable broadcast protocol [18] that can be utilized to notify all mesh routers about the new assignment. However, during configuration of the default radio, there is no guarantee that existing flows will not be disrupted. If flows can reach a destination entirely using non-default radios, those flows will not be disrupted. Flows that must use the default radio channel, on the other hand, will be disrupted until the configuration completes. However, we expect the assignment of the default radio channel to occur less frequently than for the non-default radios. This is because the interference level of the currently used default channel has to increase throughout the large deployment area of the mesh network for the average interference metric presented in Section V-B to increase enough to require re-assignment.

VII. SIMULATIONS

The objective of the simulation-based evaluation is to understand the behavior of our BFS-CA algorithm and protocol in large-scale networks. We consider four large-scale topologies. We vary the traffic patterns and the amount of external interference in three different network scenarios. In our evaluations, we compare BFS-CA against a static channel assignment. In this section, we first describe the simulation environment. We then present the four network topologies followed by a description of the three scenarios used in our evaluations. We end this section by presenting the results.

A. Simulation Environment

We use Qualnet for our simulations. We utilize the Optimized Link State Routing (OLSR) protocol [8] for routing within the mesh network and the Weighted Cumulative

Expected Transmission Time (WCETT) metric [9] for route selection. OLSR uses an optimized flooding mechanism to flood link state information in the network. WCETT is a routing metric designed for multi-radio, multi-hop wireless networks. The metric computes an estimate of the time taken to transmit a packet on a path based on the estimated bandwidth and reliability of individual links on the path and the frequency diversification of the path.

WCETT relies on the ETT metric introduced in Section VI to measure the expected delay on a path. To support the ETT metric, we implemented packet pair probing and broadcast flooding as described in Section VI-C. β in the WCETT metric can be changed to vary the weight given to the path delay and frequency diversification parameters in the metric. We set β to 0.5 which results in a good tradeoff between the two parameters during path selection [9].

OLSR link state information disseminated in the mesh is overloaded to contain link delay (ETT) values and the channel numbers on which links operate. We also modified OLSR to use source routing because it simplifies WCETT support in OLSR. We set the *TC_REDUNDANCY* protocol parameter in OLSR in order to increase the number of paths available in the mesh network.

In order to support the RFMon mode in our simulations, a radio in sniffing mode is prevented from transmitting any packets. We also simulate the layer 2 beaconing of IEEE management frames through the implementation of a beaconing module that transmits a management frame every 100 milliseconds. All radios are IEEE 802.11a radios, support 12 channels, and use auto-rate fallback. The two-ray propagation model is used with Rayleigh fading. The transmission power for the 802.11a rates is set to 18 dBm. RTS/CTS is disabled.

To make our simulations as accurate as possible, we implement interference estimation, link redirection, and the channel assignment protocol as described in Section VI. Interference estimation occurs every 5 minutes. The time spent estimating interference on each channel is set to 3 seconds. Therefore, the total time spent by a radio in RFMon mode is equal to 36 seconds (12 times 3). The CAS invokes the BFS-CA algorithm every 10 minutes. All nodes in our simulation are synchronized in time although this is not required for BFS-CA to operate correctly.

Before describing the results from our set of simulations, we first describe results from a simple topology to validate the correct operation of our algorithm. The topology is a “linear” topology consisting of four nodes. Node 1 is equipped with one radio, nodes 2 and 4 with two radios, and node 3 with three radios. Node 1 sends an 8 Mbps CBR traffic stream consisting of 1024 byte packets to node 4 that starts at 30 seconds and continues until the end of the test at 1760 seconds. Node 4 is designated as the CAS.

At the start of the simulation, all the default radios are configured to operate on a common channel. However, the non-default radios are each configured to operate on different channels. Therefore, the only path from node 1 to the gateway at the start of the simulation is on the default mesh. The first




Fig. 4. CBR Throughput in a four node network. “S” indicates periodic interference estimation by the non-default radios at each node. “A” indicates channel assignment by the CAS.

channel assignment occurs at 600 s. Just before 900 s, a lower bit rate stream is started between two interfering nodes. The interfering nodes are co-located with nodes 3 and 4 and operate on the same channel as used by the link between 3 and 4. Figure 4 shows the throughput activity at the destination in terms of number of packets received per second.

Before 600 s, node 4 receives just more than 200 packets per second from node 1. However, at 600 s, the CAS assigns the radios at nodes 2, 3, and 4 to non-overlapping channels. Consequently, the number of packets received at node 4 increased to about 1000 packets per second. At the time indicated by “S” in the figure, the non-default radios perform interference estimation. Therefore, the number of packets received at node 4 at the indicated times drops from about 1000 packets per second to a little more than 200 packets per second. Note that the number of packets delivered per second never drops to zero. This is because the nodes implement link redirection as described in Section VI-B. Just before 900 s, the interfering stream is introduced. Consequently, the number of packets received at node 4 decreases. The non-default radios at nodes 3 and 4 detect the interference at 1018 s and send their interference estimates to the CAS. At 1200 s, the CAS re-assigns the channel of the link. The figure illustrates the resulting improvement in throughput.

B. Network Topologies

We now describe the four topologies used in our large-scale evaluation. Each topology consists of 30 routers distributed in a terrain of 500x500 meters. Our choices for the number of routers and the terrain size are typical in large-scale, real-world deployments [5]. In topologies 1 and 2, the physical terrain is divided into a number of cells. Within each cell, a router is placed randomly. In generating the two topologies, we used different seed values. The two topologies reflect real-world deployments where mesh routers are uniformly distributed for maximum coverage. Topology 3 is a grid topology where the inter-router spacing is 75 meters. We consider a grid topology in order to evaluate BFS-CA in a densely populated topology. For topology 4, we choose a randomly generated topology to evaluate BFS-CA performance in an unplanned deployment of routers.

The router approximately in the center of each topology is designated as the gateway. The number of radios per router is chosen such that routers close to the gateway are equipped with more radios than ones farther away. The gateway node is equipped with four radios. Three routers, chosen from the routers one hop away from the gateway, are each also equipped with four radios. Six routers, chosen from the routers two hops away from the gateway, are each equipped with three radios. Six more routers at three hops from the gateway are equipped with two radios. Remaining routers are single-radio routers. The selection of multi-radio routers is done manually in order to comply with the router placement strategy discussed in Section III.

C. Network Scenarios

We consider three network scenarios in our evaluation. In the first scenario, we evaluate the throughput improvement obtained by utilizing multi-radio routers instead of single-radio routers. Ten randomly chosen routers at the periphery send data in two minute FTP transfers to the gateway. The scenario lasts for forty minutes. Therefore, multiple channel assignments occur in the network. This scenario is designed to be an “ideal” scenario in which: (1) there is no inter-flow interference in the network; and (2) there is no interference from external networks. To satisfy the first requirement, each source begins transmission 30 seconds after the previous one has stopped; the first source starts 620 seconds into the simulation to allow for the first channel assignment to occur. To satisfy the second requirement, we do not co-locate any external nodes.

For scenario 2, we again consider a network setting in which there is no external interference. However, unlike scenario 1, we let multiple flows within the mesh interfere with each other. The traffic is generated by the same set of sources as chosen in the first scenario. Furthermore, the sources start at the same times as in scenario 1. However, they do not stop transmitting until the end of the simulation (forty minutes).

For scenario 3, we consider a general network setting in which there is inter-flow interference as well as interference from external networks. To create this scenario, we take scenario 2 and introduce interfering nodes in multiples of 4 up to a total of 28 interfering nodes in each of the four topologies. The nodes are organized into sender-receiver pairs. Each pair is randomly distributed in the terrain space. A sender from each pair transmits a 8 Mbps CBR stream consisting of 1024 byte packets in 10 minute bursts to the receiver in the pair. The interval between each burst is 50 seconds. The nodes start transmitting at 630 s. For each 10 minute burst, the node pairs randomly select a channel for communication.

In the scenarios presented above, BFS-CA is compared against a static assignment of channels. In the static case, called the “Static Multi-Radio” scheme, the default radios operate on one channel. The first non-default radio on all multi-radio routers is tuned to a non-overlapping channel. Of the remaining routers that still have unassigned radios, the second non-default radio is tuned to a second non-overlapping

channel. Finally, the gateway and the multi-radio routers closest to the gateway tune their third non-default radio to a third non-overlapping channel.

D. Results

In scenario 1, we evaluate the throughput improvement gained by utilizing multi-radio routers instead of single-radio routers for each of the network topologies. Figure 5(a) plots the mean throughput of the ten FTP transfers in a single-radio mesh network, in a multi-radio mesh network with BFS-CA, and in a multi-radio mesh network with static assignment. The throughput improvement with multiple radios is greater than 200% for topology 1 and greater than 100% for topologies 2 and 3. The throughput improvement for topology 4, however, is only about 33% with the BFS-CA scheme and approximately 54% with the static multi-radio scheme. In topology 4, the multi-radio routers are less likely to be on paths to the gateway because they are randomly distributed in the terrain space.

More notably, Figure 5(a) indicates that in topologies 2, 3, and 4, the static multi-radio scheme performs better than BFS-CA by about 8%, 5%, and 15% respectively. This is because BFS-CA improves throughput by tuning individual links to non-overlapping frequencies. Because of this strategy, the opportunity to find channel-diversified paths is less than with the static scheme where regions of the mesh network are tuned to non-overlapping frequencies. Therefore, with the static scheme, the opportunity to choose channel-diversified paths is better. As a result, the static scheme performs better.

In order to verify the above reasoning, we computed the Channel Diversity Extent (CDE) of all flows with the two schemes. The CDE of a flow is obtained by taking the average of the CDEs of all paths traversed for the flow; a flow can traverse multiple paths because of variations in a link’s ETT value, which can result in path changes. The CDE of a path is defined as the ratio of the number of channels used in the path to its hop-count. For example, if a 4 hop path makes use of 4 channels, its CDE is 1. On the other hand, if the path makes use of only 1 channel, its CDE is 0.25. A path with a high CDE is generally preferred over a path with a low CDE. A TCP flow is bi-directional due to acknowledgments and is considered as two flows in our calculations.

Figure 5(b) is a scatter plot of the CDEs of all flows for the two multi-radio schemes. The x value of each point in the plot is the CDE of a path with BFS-CA, and the y value is the CDE of the same path with the static scheme. The line $x = y$ indicates paths with equal CDE values. A majority of the points in the plot are above the $x = y$ line. These points indicate paths that with the static scheme have higher CDE values than with the BFS-CA scheme. Consequently, paths with the static scheme are more channel diversified than with the BFS-CA scheme.

Figure 5(c) is a plot of the average hop-counts for the two multi-radio schemes in scenario 1 for the four topologies. The hop-count with BFS-CA is higher than with the static scheme. This is because WCETT, the routing metric in our evaluations,




Fig. 5. Results from scenario 1 where the BFS-CA and static schemes are compared against a single-radio scheme. Scenario 1 has no inter-flow interference and no external interference.




Fig. 6. Results from scenario 2 where BFS-CA is compared against the static scheme in the presence of inter-flow interference but without external interference.

prefers a frequency diversified path over alternatives that may be shorter. Because of the assignment strategies of the two schemes, WCETT picks longer frequency diversified paths with BFS-CA than with the static scheme.

The above noted performance improvement of the static scheme over BFS-CA is achieved only in an ideal scenario where sources send traffic to the gateway one at a time. Because of the assignment strategy in the static scheme, when multiple sources transmit, flows to the gateway each traverse links tuned to the same channels, resulting in increased interference between the flows. Consequently, the throughput improvement with the static scheme is expected to be less than with the BFS-CA scheme.

This intuition is verified in scenario 2 wherein multiple sources transmit to the gateway at the same time. Figure 6(a) plots the mean throughput obtained with the two multi-radio schemes. Clearly, BFS-CA outperforms the static scheme. Specifically, the throughput improvement is as high as 72% with topology 2 and 48%, 60%, and 13% with topologies 1, 3, and 4, respectively. Figure 6(b) shows the CDEs of the flows for the two schemes. Points lie in approximately equal numbers above and below the $x = y$ line. Although the flows in the two schemes are equally channel diversified, the throughput with BFS-CA is higher because of reduced inter-flow interference. Figure 6(c) plots the average hop-count with the two schemes. The average hop-count with BFS-CA is slightly higher than with the static scheme. This is because of the assignment strategy used by BFS-CA. However, with

BFS-CA, the flows traverse paths that result in less inter-flow interference than with the static scheme. Therefore, the mean throughput with BFS-CA is greater.

In the third scenario, we consider the effect of varying amounts of interference on the performance of the mesh topologies for the two multi-radio schemes. Figure 7 shows the percentage difference in the mean throughput of the ten FTP transfers achieved by the two schemes. In a majority of the cases, BFS-CA outperforms the static scheme. As an average of the performance differences with varying the number of interfering radios, BFS-CA performs better than the static scheme by 42.14%, 31.14%, 15.75%, and 11.85% for topologies 1, 2, 3, and 4 respectively. BFS-CA's improved performance over the static scheme occurs because the mesh routers in the network topologies are able to detect the increased interference and are therefore re-assigned by the CAS to channels different than ones used by the interfering nodes. In the case of topology 4, the unplanned distribution of mesh routers in the terrain space results in BFS-CA performing only slightly better than the static scheme. More noteworthy is BFS-CA's performance in topology 3 (the grid topology) in which the throughput gains with BFS-CA is only marginally better than the topology 4 case. This is because in topology 3, the mesh routers and interfering networks are densely populated in the terrain space. Consequently, the interference foot-print of the interfering networks is greater in the dense environment. As a result, re-assignment of channels yields only minimal throughput improvement with BFS-CA.

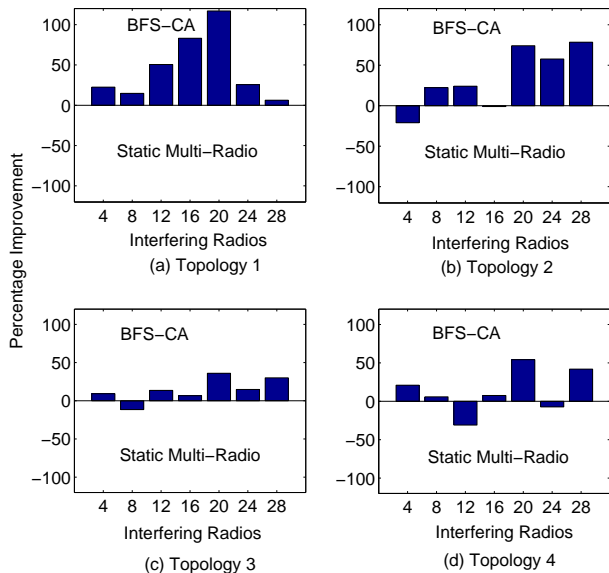


Fig. 7. Percentage improvement in throughput with BFS-CA relative to Static Multi-Radio in the presence of interference. The number of interfering radios is shown on the x-axes.

VIII. PROTOTYPE IMPLEMENTATION

Our primary motivation in implementing a prototype is to demonstrate the practicality of our proposed BFS-CA algorithm and protocol. We evaluate the implementation in a six node IEEE 802.11b testbed with the Linux 2.4.26 operating system.

The prototype implementation consists of a router module and a CAS module. The router module is installed on each mesh router and implements the interference estimation and link redirection procedures as described in Section VI. To perform RFMon sniffing, it invokes Linux utilities such as the *iwconfig* wireless utilities and the *tethereal* packet capture tool. The CAS module implements the BFS-CA algorithm. Interference estimation occurs every 5 minutes, and the channel selection procedure is invoked every 10 minutes as in the simulations.

Because IEEE 802.11b supports only three non-overlapping channels (1, 6, and 11), the performance improvement obtained using the proposed solution in our testbed setting is expected to be limited. We considered using Atheros based IEEE 802.11a radios. However, our testbed nodes use Linux and *ad-hoc mode* support in Linux drivers for IEEE 802.11a radios is currently faulty [3].

All nodes in the topology are IBM R32 Thinkpad laptops. The six nodes are each equipped with one Netgate NL-2511CD PCMCIA radio. Nodes B and C are each equipped with an additional Linksys WUSB12 USB radio. We use USB extender cables for the Linksys radios. Nodes A, B, C, and D form a static multi-hop wireless network in a simple line topology, i.e., A cannot see C and D, B cannot see D, and vice-versa. Nodes E and F are interfering nodes positioned roughly a meter away from B and C respectively. Nodes E

and F use a different ESSID and network IP address than A, B, C, and D. All radios are configured to operate at 11 Mbps. RTS/CTS is disabled.

For the entire duration of the experiment, the default mesh is set to channel 1. We did not allow the re-assignment of the default mesh channel during our experiments because our WUSB12 radios do not support dynamic reconfiguration in ad-hoc mode. Node B is designated as the CAS. The CAS informs nodes A, C, and D about the IP address range that belongs to the mesh network. At the start of the experiment, the radios on E and F and the non-default radios on B and C are tuned to channel 6.

In the experiment, A sends 1024 byte UDP packets as fast as possible to saturate the path to D for a duration of 40 minutes. The interfering node E sends 1024 byte UDP packets as fast as possible to saturate the link to node F starting at 250 s until the end of the experiment. At 1400 s, we manually switch nodes E and F from channel 6 to channel 11 in order to measure the response of our algorithm implementation to a varying interference.

Figure 8 illustrates the throughput in terms of number of packets received at D and F. From the figure, it can be clearly seen that the throughput of the stream from A to D is *not affected* by the interfering stream. The reason is that the links A-B and C-D are in interfering range of each other. Consequently, these links form the bottleneck. As a result, link B-C is able to consistently sustain the low number of packets injected by A in spite of interference from link E-F. The drop in throughput every 300 seconds is because of the periodic interference estimation performed by the non-default radio on B and C. The estimation occurs for 9 seconds (3 seconds times 3 non-overlapping channels). The throughput never drops to zero because link redirection is performed over the default radios until the interference estimation completes. At the times indicated by *S* in the figure, the non-default radios at B and C detect the interference and notify the CAS. At the times indicated by *A*, the non-default radios are re-assigned to a channel not used by the interfering stream. The result of the re-assignment can be clearly seen in the throughput plot of the interfering network. After the first re-assignment, the link B-C switches from Channel 6 to 11. Consequently, the throughput of the interfering stream increases. At 1400 s, link E-F is manually assigned to channel 11. As a result, the throughput drops. It increases again when link B-C switches from Channel 11 back to Channel 6 due to the second re-assignment.

IX. RELATED WORK

There exists a large number of studies that aim to address the capacity problem in wireless mesh networks. We summarize a representative sample below.

Several proposals focus on improving the IEEE 802.11 MAC protocol to support multiple channels [7], [11], [23]. The key advantage of such schemes is that only a single radio is required to support multiple channels. The disadvantage, however, is that they require changes to the MAC layer and the




Fig. 8. Throughput in the presence of an interfering stream.

hardware in order to support per-packet channel switching. To the best of our knowledge, such hardware is still not available.

Raniwala et al. [19], [20] propose a load-aware algorithm to dynamically assign channels in a multi-radio mesh network. Raniwala's proposal requires that anticipated traffic loads and paths traversed by flows be known before channel assignment takes place. The key distinguishing aspect in our proposal is that we assign channels based entirely on knowledge of interference in the mesh network. Kyasanur et al. propose a hybrid channel assignment solution for multi-radio wireless ad hoc networks [16]. In this scheme, a subset of radios on a router are statically assigned channels. The remaining radios dynamically switch to the static channels assigned to neighboring routers in order to communicate with them. The scheme requires that radios can switch between channels on a per-packet basis. To the best of our knowledge, such radios are not available in the marketplace.

Some mesh hardware vendors [2], [4] offer multi-radio mesh routers that utilize proprietary channel assignment schemes. Therefore, we are unable to provide a fair comparison.

X. CONCLUSIONS

Multi-radio routers can significantly improve the performance of wireless mesh networks. However, any static assignment of channels to the mesh radios can degrade network performance because of interference from co-located wireless networks. This paper presented BFS-CA, a dynamic, interference-aware channel assignment algorithm and corresponding protocol for multi-radio wireless mesh networks. BFS-CA improves the performance of wireless mesh networks by minimizing interference between routers in the mesh network and between the mesh network and co-located wireless networks. The proposed solution is practical and easily implementable. We find that BFS-CA results in significant performance improvements in the presence of varying interference levels, which are validated through empirical measurements on a testbed. As future work, we plan to evaluate BFS-CA on the UCSB MeshNet [6], a thirty node multi-radio wireless mesh testbed at UCSB.

REFERENCES

[1] Aeroflex 6970 RF Power Meters. <http://www.aeroflex.com>.
 [2] BelAir Networks. <http://www.belairnetworks.com>.

[3] MADWIFI Project. <http://sourceforge.net/projects/madwifi/>.
 [4] MeshDynamics. <http://www.meshdynamics.com>.
 [5] Metro-Scale WiFi As City Service chaska.net, Chaska, Minnesota. http://www.tropos.com/pdf/chaska_casestudy.pdf.
 [6] UCSB MeshNet Project. <http://moment.cs.ucsb.edu/meshnet>.
 [7] P. Bahl, R. Chandra, and J. Dunagan. SSCH: Slotted Seeded Channel Hopping For Capacity Improvement in IEEE 802.11 Ad Hoc Wireless Networks. In *ACM MobiCom*, Philadelphia, PA, September 2004.
 [8] T. Clausen and P. Jacquet. Optimized Link State Routing Protocol. Internet Engineering Task Force, RFC 3626, October 2003.
 [9] R. Draves, J. Padhye, and B. Zill. Routing in Multi-radio, Multi-hop Wireless Mesh Networks. In *ACM MobiCom*, Philadelphia, PA, September 2004.
 [10] P. Gupta and P. Kumar. Capacity of Wireless Networks. In *IEEE Transactions on Information Theory*, volume 46, pages 388–404, March 2000.
 [11] N. Jain, S. Das, and A. Nasipuri. A Multichannel CSMA MAC Protocol with Receiver-Based Channel Selection for Multihop Wireless Networks. In *IEEE International Conference on Computer Communications and Networks*, Scottsdale, AZ, October 2001.
 [12] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the Real-World Performance of Carrier Sense. In *ACM Sigcomm Workshop on Experimental Approaches to Wireless Network Design and Analysis*, Philadelphia, PA, August 2005.
 [13] A. Jardosh, K. Ramachandran, K. Almeroth, and E. Belding. Understanding Congestion in IEEE 802.11b Wireless Networks. In *ACM/USENIX International Measurement Conference*, Berkeley, CA, October 2005.
 [14] J. Katzela and M. Naghshineh. Channel Assignment Schemes for Cellular Mobile Telecommunications Systems: A Comprehensive Survey. In *IEEE Personal Communications*, volume 3, pages 10–31, June 1996.
 [15] S. Keshav. A Control-Theoretic Approach to Flow Control. In *ACM Sigcomm*, Zurich, Switzerland, September 1991.
 [16] P. Kyasanur and N. Vaidya. Routing and Interface Assignment in Multi-Channel Multi-Interface Wireless Networks. In *IEEE Wireless Communications and Networking Conference*, New Orleans, LA, March 2005.
 [17] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of Link Interference in Static Multi-hop Wireless Networks. In *ACM/USENIX International Measurement Conference*, Berkeley, CA, October 2005.
 [18] E. Pagani and G. Rossi. Reliable Broadcast in Mobile Multihop Packet Networks. In *ACM MobiCom*, Budapest, Hungary, September 1997.
 [19] A. Raniwala and T. Chiueh. Architecture and Algorithms for an IEEE 802.11-based Multi-Channel Wireless Mesh Network. In *IEEE Infocom*, Miami, FL, March 2005.
 [20] A. Raniwala, K. Gopalan, and T. Chiueh. Centralized Channel Assignment and Routing Algorithms for Multi-Channel Wireless Mesh Networks. In *Mobile Computing and Communications Review*, volume 8, pages 50–65, April 2004.
 [21] K. Rosen. *Discrete Mathematics and its Applications*. McGraw Hill, 1999.
 [22] K. Sanzgiri, I. Chakeres, and E. Belding-Royer. The Utility of Perceptive Communication between Distant Wireless Nodes. In *IEEE/Create-Net TridentCom*, Barcelona, Spain, March 2006.
 [23] J. So and N. H. Vaidya. Multi-Channel MAC for Ad Hoc Networks: Handling Multi-Channel Hidden Terminals Using a Single Transceiver. In *ACM MobiHoc*, Tokyo, Japan, May 2004.
 [24] A. Vasani, R. Ramjee, and T. Woo. ECHOS - Enhanced Capacity 802.11 Hotspots. In *IEEE Infocom*, Miami, FL, March 2005.
 [25] S. Vasudevan, J. Kurose, and D. Towsley. Design and Analysis of a Leader Election Algorithm for Mobile Ad Hoc Networks. In *IEEE International Conference on Network Protocols*, Berlin, Germany, October 2004.