Traffic-Aware Channel Width Adaptation in Long-Distance 802.11 Mesh Networks

Sofia Pediaditaki[*]
Intel Research Center
Intel Labs Barcelona – UPC

Mahesh K. Marina and Daniel Tyrode
School of Informatics
University of Edinburgh

ABSTRACT

We consider the traffic adaptive channel allocation problem in long-distance 802.11 mesh networks. Our approach is to exploit the capability provided by 802.11 hardware to use different channel widths and assign channel widths to links based on their relative traffic volume. We show that this traffic-aware channel width assignment problem is NP-complete and propose a polynomial time, greedy channel allocation algorithm that guarantees valid channel allocations for each node. Evaluation of the proposed algorithm via simulations of real network topologies shows that it consistently outperforms the current approach of fixed width allocation due to its ability to adapt to spatio-temporal variations in traffic demands.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*

General Terms

Algorithms, Design, Performance

Keywords

Long-distance 802.11 mesh networks, Channel allocation, Traffic awareness, Channel width assignment, Adaptive spectrum management, Resource allocation

1. INTRODUCTION

The remarkable success of WiFi (based on the IEEE 802.11 standard) has led to its use in originally unintended scenarios. Long-distance 802.11 mesh network scenario, the focus of this paper, is one of those that is making a huge impact in the real world in helping bringing low cost Internet access to rural areas and developing regions (e.g., [1, 2]) by enabling affected communities and new Wireless Internet Service Providers (WISPs). The main impediment for provisioning broadband access in these regions is the

---

[*]This work was done when this author was at the University of Edinburgh.
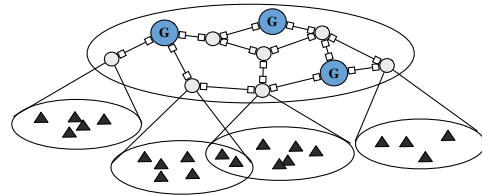
Figure 1: Multi-tier 802.11 based mesh network model.

deployment cost to serve either low density scattered communities or populations with limited incomes. It will take years to fully penetrate into these underserved regions offering limited profitability with broadband access technologies prevalent in urban areas (fiber, DSL, cable, 3G/4G) because of their high infrastructure costs. The long-distance 802.11 mesh scenario is also of interest as part of larger and scalable mesh networks in urban settings along the lines of the architecture considered in [3]. It also fits well with the needs of WISPs that span both rural and urban areas.

We consider a tiered 802.11 based mesh network model shown in Fig. 1 applicable to rural and urban settings mentioned above. Two tiers are shown of which *our focus is on the topmost "directional backhaul" tier*. Nodes in the top tier could be separated potentially by long distances in the order of several Kms, hence their interconnection into a network is achieved with the use of a pair of high-gain directional antennas per link. As such this tier can be seen as a point-to-point wireless network. Some of these nodes in the top tier called gateway nodes connect to the wired Internet infrastructure (e.g., nodes labeled 'G' in Fig. 1). While some nodes in the top tier only have the router role to forward data between other top tier nodes, several nodes additionally provide connectivity to the lower tier subnets using point-to-multipoint wireless links as illustrated. Therefore the latter set of nodes can be seen as traffic aggregation points in the directional backhaul tier. Each of the subnets in the lower tier could in turn be an omnidirectional mesh network with each node representing a rooftop mesh access point (in a village or urban neighborhood). One can imagine an additional tier (not shown in the figure) connecting devices inside homes to a rooftop access point.

The directional backhaul tier in the network model is a specific type of multi-radio multi-channel mesh network. Each node in the backhaul has as many radio interfaces as the number of incident links (each connected to a directional antenna), and each of these links are assigned a different channel to avoid *side-lobe interference* that occurs with commonly used high-gain directional antennas [4] — non-negligible side-lobe energy from directional transmission on a link appears as interference to reception on other

co-incident links, so such interference needs to be avoided. More-over, for long-distance communication, besides directional antennas, higher radio transmit power may also be needed. Therefore, such long-distance point-to-point wireless communication is restricted by spectrum regulatory bodies to a few specified frequency bands with relatively higher transmit limits. The 5.8GHz frequency band is one such band and is available in most regions of the world. Consequently, the total amount of spectrum available for the directional backhaul tier is limited (e.g., 100MHz in the 5.8GHz band as opposed to more than 500MHz available for indoor wireless LAN operating on the 5GHz unlicensed bands).

Since the directional backhaul tier serves as an intermediate data transport network between the wired Internet and large number of client devices in the lower tiers, limited available spectrum needs to be managed judiciously and adaptively in response to varying traffic demands. Long-distance mesh deployments in practice tend to skirt around this important issue for lack of a suitable adaptive channel allocation framework. In fact, it is common to assign identically sized but possibly different channels to network links at deployment time and have them remain unchanged (e.g., [5]) or even use only a single channel for the whole network (e.g., [4]).

In this paper, we aim to fill this void by viewing *channel width* as a knob to enable traffic-aware channel allocation in long-distance 802.11 based mesh networks, i.e., focusing on the directional backhaul tier in Fig. 1. The intuition behind our approach is as follows: since the total amount of available spectrum is limited, adapting to spatio-temporal variations in traffic demand can be achieved by allocating "wider" channels to links with higher demand by taking spectrum away from links with less demand. In other words, greater capacity is assigned to heavily utilized links, thereby benefiting the flows passing through them. Following the work of Chandra et al. [6] who first demonstrated experimentally the throughput, range and energy efficiency benefits of channel width adaptation in an isolated 802.11 link scenario, other research efforts have since highlighted the value of channel width adjustment in 802.11 wireless LANs [7, 8]. To the best of our knowledge, we are not aware of work that considers channel width adaptation in the context of long-distance mesh networks.

Our contributions are as follows: (1) We present a graph theoretic formulation of the traffic-aware channel width assignment problem in long-distance mesh networks and show that it is NP-complete (Section 3). (2) We develop a polynomial time algorithm for assigning channel widths to links based on their relative traffic volume; the algorithm ensures that every node gets a valid channel allocation (Section 4). (3) Our simulation based evaluation of the proposed algorithm using real network topologies shows that it delivers substantial improvements in performance (40-70% throughput improvements) from adapting the channel allocation in response to variations in spatio-temporal traffic demands (Section 5).

We start by looking at related work in a bit more detail in the following section.

2. RELATED WORK

Related work falls into three categories as discussed below.

2.1 Channel Allocation and MAC Design for Long-Distance Mesh Networks

Most of the work in this space focuses on TDMA-based MAC protocols as an alternative to 802.11. While part of the motivation behind these protocols is the detrimental impact of high propagation delays on 802.11 performance for *very* long distance wireless links in the order of 100Kms, the rest has to do with the so-

called "side-lobe interference" issue [4]. The latter refers to the interference among incident (directional) links at a node using the *same* channel, especially when one or more of them are transmitting and other links are receiving. This type of interference occurs with commonly used high gain directional antennas having non-negligible side lobes in their radiation pattern. The 2P protocol [4] is the first alternative design in the literature to address the above problems using a TDMA based approach that requires each node to alternate between transmitting (on all incident links) and receiving (again, on all incident links). Several subsequently proposed channel allocation protocols assume 2P as the underlying MAC protocol (e.g., [9, 10]). However, the 2P protocol works only if the network topology is a bipartite graph. This limitation has been addressed in a later proposal called JazzyMAC [11]. By their very nature, both 2P and JazzyMAC need inter-nodal time synchronization, an additional requirement.

Our work on channel allocation in this paper instead assumes standard 802.11 MAC, based on the following two observations: (1) Real world long-distance wireless links are typically in the order of several Kms to few tens of Kms for which 802.11 MAC gives acceptable performance through suitable adjustment of the built-in ACK timeout. This is experimentally shown in [12] and is also confirmed by our experience deploying and monitoring the Tegola network in rural Scotland [13] for the past three years with links in the range of 2-20Km. (2) The use of multiple channels alleviates the side-lobe interference problem as different links at a node can be assigned to different non-interfering channels.

The recent channel assignment work of Dutta et al. [14] is similar in spirit to ours in that it also assumes standard 802.11 MAC, thereby leveraging readily available commodity 802.11 hardware. However, that work focuses on *directed* edge coloring, requiring a channel for each *directed* link, potentially resulting in inefficient spectrum utilization with limited spectrum; it also increases the cost and deployment complexity — to support directed edge coloring, each node requires two directional antennas (and radios) per link and those antennas need to be carefully separated to manage side lobe interference.

Crucially, to the best of our knowledge, the existing channel allocation literature on long-distance mesh networks does not consider channel width adaptation, which is the key aspect of our work.

2.2 Channel Width Adaptation in 802.11 Networks

Chandra et al. [6] were the first to examine the impact of channel width adaptation on throughput, range and power consumption and obtain experimental evidence of the potential benefits of adjusting the width of a channel in 802.11 networks. Their focus, however, is on the simplest case, i.e., single link, for which they propose a channel width adaptation algorithm called SampleWidth. Most of the subsequent work on channel width adaptation in 802.11 networks has focussed on the wireless LAN (WLAN) scenario. Moscibroda et al. [7] consider channel width adaptation for achieving load balancing in multi-AP WLANs as an alternative approach to transmit power and client association controls. They show that the dynamic width assignment problem is NP-hard and propose heuristics that are evaluated using simulations. Yuan et al. [15] take a game-theoretic approach to address the same problem and propose a decentralized learning-based algorithm for achieving optimal allocation. Recently, Rayanchu et al. [8] find that realizing the benefits of variable channel widths in multi-AP 802.11 based WLANs requires reliable characterization of interference (conflicts) and they propose a mechanism to efficiently model such conflicts.

Note that adapting these WLAN channel width adaptation pro-

posals to our multihop wireless network context is not straightforward because of the fundamental differences between the two scenarios. For instance, we have the requirement to maintain network connectivity *wirelessly*, whereas access points (APs) in a wireless LAN are interconnected via a wired backhaul network.

In [16], the authors show that splitting the spectrum into variable width channels among mutually interfering transmitters and keeping them active simultaneously is more effective than having them share a larger, fixed-width channel because of SNR increase at smaller channel widths for a given transmit power. A similar argument can be used in our scenario for determining the feasible combinations of channel widths and bit-rates (modulation and coding schemes) for an installed long-distance link with fixed transmit power and antenna gains. We also note that Orthogonal Frequency Division Multiple Access (OFDMA) used in WiMAX [17], a fine-grained approach operating at the subcarrier level for spectrum sharing among multiple users, is related to channel width adaptation in 802.11 based networks.

2.3 Channel Allocation in Multi-Radio Multi-Channel Wireless Mesh Networks

Much work has been done on channel allocation in multi-radio multi-channel wireless mesh networks using omnidirectional antennas (e.g., [18, 19, 20]). Though the network model we consider is a special type of a multi-radio multi-channel wireless mesh network, our interference model is different from mesh networks based on omnidirectional antennas; more importantly, we have an added decision variable — channel width — that is not present in the above works. Recently a few papers looked into this aspect (e.g., [21, 22]), focusing on the omnidirectional mesh network scenario and mainly employing mathematical optimization methods (e.g., mixed integer linear programming). More recently, Wu et al. [23] consider the adaptive width channel allocation in multi-radio wireless mesh networks from a game theoretic perspective but their work is associated with unrealistic assumptions such as all nodes lying within a single collision domain, and each node participating in only one communication session and that too over a single hop.

In DMesh [24], the authors propose a wireless mesh network architecture based on directional antennas with the goal of exploiting spatial reuse benefit of directional communication to go along with the benefit of using multiple channels. We, on the other hand, are more interested in the range benefit offered by directional antennas.

3. MODEL AND PROBLEM FORMULATION

As stated at the outset, we consider a multi-tier 802.11 mesh network scenario as shown in Fig. 1 and our focus in this paper is on the topmost directional backhaul tier, especially keeping the rural wireless Internet access use case in mind. We model the network topology of the directional backhaul tier as an undirected graph, $\mathcal{T} = (\mathcal{N}, \mathcal{L})$. Each node n at the backhaul tier is equipped with $K_n (K_n \geq 1)$ 802.11 wireless interface cards, each attached to a directional antenna forming an end of a long-distance point-to-point wireless link with a neighboring node. We use the notation $n_p, 1 \leq p \leq K_n$ to refer to the p^{th} interface at n. Note that K_n is equal to the number of point-to-point wireless links incident at node n.

We use (n, m) to denote the logical point-to-point link between two nodes n and m. And we use (n_p, m_q) to denote the actual physical bidirectional point-to-point link between the p^{th} interface of node n and the q^{th} interface of neighbor m. When we need to refer to the direction of a link, we will use the notation $n_p \rightarrow$
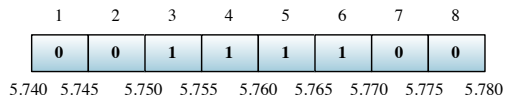
Figure 2: An example to illustrate *Spectrum Allocation Map (SAM) of a node*. In this example, the total spectrum available is assumed to be 40MHz, which results in 8 5MHz wide blocks. Each 5MHz wide spectrum block is shown using a box with a block identifier above the box and start frequency of the block underneath. Based on SAM of the node (i.e., the values shown inside boxes), blocks 3-6 are used while the rest are free. Assuming the node has only one link and all the used blocks correspond to the channel allocated to that link, the *block assignment (channel) for the link* is: $ch = <f_s, w> = <5.75GHz, 20MHz>$ with center frequency, $f_c = 5.76GHz$.

m_q to refer to the direction from n to m. Note that, in practice, each link in a long-distance 802.11 mesh network is determined at deployment time by pointing a pair of directional antennas located at two mast sites towards each other. These links remain fixed when the network goes into operational stage; our focus in this paper is solely on adapting the "channel" used by a link based on the traffic volume measured over it.

We use $\mathcal{F}$ to denote the total available spectrum. In the case of widely used 5.8GHz band, $\mathcal{F} = 100$MHz, ranging from 5.725GHz to 5.85GHz with small guard bands on either end[1]. This band can accommodate 5 20MHz channels (the default channel width in 802.11a) — channel numbers: 149, 153, 157, 161 and 165. The same band can be used to accommodate up to 2 40MHz channels. We assume the set of available channel widths to be 5, 10, 20, 40 based on what is currently supported by commodity 802.11 hardware.

A channel in our context is defined by the tuple $ch = <f_c, w>$, where f_c represents the center frequency and w the width of the channel taking one of the 4 values just mentioned — frequencies of ch range from $f_c - w/2$ to $f_c + w/2$. For example, 802.11a channel number 149 corresponds to a 20MHz channel centered at 5.745GHz. Given $\mathcal{F}$, the range of frequencies that fall within the spectrum and available channel widths, several channels can be realized by choosing various center frequencies and widths. We can alternatively look at ch using its start frequency, $f_s = f_c - w/2$ and width w; in this case, ch ranges from f_s to $f_s + w$. We use the notation $ch_{(n_p, m_q)}$ to refer to the channel assigned to the link (n_p, m_q). Note that $ch_{n_p \rightarrow m_q} = ch_{m_q \rightarrow n_p}$.

For convenience, we view the given spectrum as a sequence of atomic 5MHz wide *blocks*. For example, when the available spectrum is 100MHz, we have 20 5MHz wide blocks. If S denotes the number of blocks[2] (20 in the example), then we assume that $S \geq 2 * \Delta(\mathcal{T}) - 1$, where $\Delta(\mathcal{T})$ is the maximum node degree in $\mathcal{T}$. This is quite a reasonable assumption since a node in the backhaul directional tier typically has at most around a handful of incident point-to-point wireless links. The relevance of this assumption will become clear later on in Section 4.

We now define a *spectrum allocation map for each node n: SAM_n* that represents the spectrum usage of the interfaces at n. Specifically, this map is a sequence of bits associated with the blocks, where a bit is set to 1 if the corresponding block is occupied by some interface of the node. Otherwise, it is 0. Fig. 2 shows

---

[1] 10MHz at the lower end and 15MHz on the upper end.
[2] Henceforth, we just use the term 'block' as a shorthand for '5MHz wide block'.

an example. Using the notion of blocks, the channel assigned to a link (n_p, m_q), $ch_{(n_p,m_q)}$, can be seen as contiguous and identical set of blocks in SAM_n and SAM_m; we refer to such assignment of contiguous set of blocks to a link (n_p, m_q) as the *block assignment for the link*, denoted by BA_{n_p,m_q} (see Fig. 2 for an example).

Having described what a channel means in our model, we now introduce the three key constraints in our channel allocation problem.

Side-lobe interference constraint:
$BA_{n_p,m_q} \cap BA_{n_r,l_s} = \emptyset; n, m, l \in \mathcal{N}, \exists (n_p, m_q), \exists (n_r, l_s),$
$p \neq r, m \neq l.$ (**C1**)

This constraint essentially requires that any two incident links at a node are assigned different *non-overlapping* channels[3].

Minimum channel constraint:
$ch_{n_p,m_q} = <f_c, w>$ s. t. $w \geq 5; \exists (n_p, m_q).$ (**C2**)

This constraint requires that each link is assigned at least a 5MHz wide channel (i.e., a block). This is to make sure that network topology always remains intact with each link having a usable channel, which can be used for exchanging at least control traffic (e.g., routing messages).

Total spectrum constraint:
$0 < \sum_{m, \exists (n_p,m_q)} ch_{n_p,m_q} \leq \mathcal{F}.$ (**C3**)

This constraint makes sure that channels allocated to incident links at a node do not exceed the total spectrum available.

Before going to the objective function, we need to model flow (traffic) on a link and the link capacity based on the channel allocated to it. Let f_{n_p,m_q} denote the total traffic flow passing through a link (n_p, m_q): $f_{n_p,m_q}{}^4 = f_{n_p \to m_q} + f_{m_q \to n_p}$. The capacity of a link (n_p, m_q) denoted by C_{n_p,m_q} is dependent on the channel it is allocated and link characteristics (link distance, etc.). Channel width and best bit-rate (modulation and coding scheme) supported by the link are inter-dependent, and together determine the raw link capacity, $C_{n_p,m_q}{}^5$.

Objective function: our objective can be stated as *minimizing the maximum excess link load across all links in the network* subject to constraints **C1**-**C3** at each node, where excess link load of a link (n_p, m_q) is defined as $max(f_{n_p,m_q} - \delta * C_{n_p,m_q}, 0)$. Here δ models the fraction of raw link capacity effectively available at network

---

[3]Here we make the simplifying assumption that side lobe interference can be avoided if co-incident links are assigned channels with frequency ranges that do not overlap. We can extend it to incorporate a sophisticated non-binary model of interference between any two co-incident links based on the work by Angelakis et al. [25] that takes into account antenna radiation patterns, inter-antenna distance, separation between center frequencies of channels assigned to the two links and channel widths. We elaborate more on the latter in Section 6.

[4]A lightweight sampling method to continually estimate the total traffic flow on a link based on MRTG and SNMP is described in [26]; Fig. 7 is generated from data obtained using this method in the Tegola network [13]

[5]Prior work [27, 28] has shown that typical rural long-distance wireless links, the particular focus of our work, experience negligible link quality variations. In such cases, we only need to consider interaction between channel width and bit-rate to estimate link capacity based on [6] for a given link quality (that can be obtained via measurement) . As a corollary, bit-rate would also be stable in the fixed width case.

layer after discounting overhead related to link and physical layers (e.g., headers, inter-frame spaces) as well as routing control traffic overhead. For example, if 20MHz channel width can support 54Mbps physical layer bit-rate, then effective achievable capacity above the link layer is at most 30Mbps dependent on frame length. Assuming some portion of it (say 10%) is consumed by control traffic, the maximum effective capacity is 27Mbps, leading to a δ value of 0.5.

The intuition behind using this objective function is to evenly distribute the available spectrum resource among links based on their traffic demands so that capacity assigned to a link matches its load as closely as possible.

We refer to the decision problem equivalent of the above optimization problem as *Channel Width Assignment for MinMax Excess Link Load*, which can be stated as follows: Is there a channel width assignment such that the maximum excess link load $\leq$ B (where B is a non-negative integer)?

THEOREM 1. *The* Channel Width Assignment for MinMax Excess Link Load *decision problem as stated above is NP-complete.*

PROOF. A channel width assignment can be verified in polynomial time, thus the problem is clearly NP.

The rest of the proof is by restriction. We show that the above channel width assignment problem contains the minimum edge coloring problem (also called the minimum chromatic index) - a known NP-complete problem - as a special case [29].

Specifically, we show that a specific instance of the problem at hand is identical to the minimum edge coloring problem. For this instance, the following constraints hold: (i) only one channel width is allowed (5MHz), (ii) the total amount of spectrum available = maximum node degree * 5MHz, (iii) $f_{n_p,m_q} << \delta * C_{n_p,m_q}$, for all links (n_p, m_q) when using a 5MHz wide channel, and (iv) the graph is chosen to be identical for both problems and bound for number of colors (chromatic index) for the edge coloring is set to the maximum node degree.

The above specified instance of our problem is identical to the minimum edge coloring problem, which completes the proof. $\square$

Note that using the above stated objective function and per-node constraints **C1**-**C3**, we can formulate this problem as a mixed integer linear program for obtaining a lower bound on the optimum.

4. CHANNEL WIDTH ASSIGNMENT ALGORITHM

In this section, we describe a polynomial time greedy heuristic for traffic-aware channel width assignment in long distance 802.11 mesh networks.

In our approach, channel allocation is performed independently from routing but influences it and vice versa. Moreover, channel reallocation is done at a relatively slower timescale than routing. This is because channel width changes are more disruptive as each such change at least requires endpoints of a link to reconfigure (and even restart) their corresponding wireless interfaces. Exact frequency of channel width adaptation is a tradeoff between responsiveness to traffic dynamics and keeping network disruption and overhead low; we will discuss this issue further in Section 6. Keeping routing and channel allocation independent has the advantage that any routing protocol can be used on top of the channel width adaptation algorithm. While a traffic adaptive routing protocol such as [30] would be an ideal companion for the channel allocation algorithm, some network deployments may not plan for enough redundancy in the
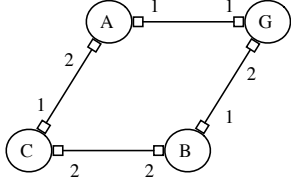
Figure 3: An example network.

	Interface 1	Interface 2
(1)	20	20
(2)	20	10
(3)	20	5
(4)	10	20
(5)	10	10
(6)	10	5
(7)	5	20
(8)	5	10
(9)	5	5

Figure 4: Feasible channel width combinations for node A in Fig. 3 under the assumption that A_1 has a higher priority than A_2.

	Interface 1	Interface 2
(1)	20	20
(2)	20	10
(3)	20	5

Figure 5: Valid channel width combinations for node A in Fig. 3 after its interface 1 is colored by node G with a 20MHz channel.

topology for cost reasons, making the complexity of using a traffic adaptive routing protocol questionable; in such cases, a single path routing protocol can be used with the burden of traffic adaptivity shifted largely to the channel allocation algorithm. For our simulation based evaluations, we have implemented P-STARA [30] and use it as the default routing protocol.

Before describing the algorithm, we introduce the concepts of *feasible* and *valid* channel allocations. A feasible channel allocation for a node n is an allocation that respects constraints **C1**-**C3** (see Section 3). As an example, consider node A in the example network shown in Fig. 3 and suppose that the total amount of available spectrum is 40MHz. Also suppose that interface 1 (link to G) at A has a higher priority than interface 2 at A (link to C) — setting of node and link priorities is discussed shortly. Under these assumptions, feasible channel width combinations for node A are shown in Fig. 4 in the order of decreasing preference[6]. Not all of these combinations are valid depending on when node A gets to do its channel allocation with respect to its neighboring nodes. To see this, consider the network in Fig. 3 again. Suppose that node G allocates a 20MHz channel to its interface numbered 1, thus coloring[7] link to A and corresponding interface at A. This coloring decision by G immediately reduces the possible width combinations to the set shown in Fig. 5. Moreover, depending on where channel for link (G_1, A_1) lies in the spectrum, the possible width combinations may reduce even further. If the SAM_A after link (G_1, A_1) is allocated 20MHz channel looks like in Fig. 2, then the combination (1) in Fig. 5 is no longer possible. In the worst case, for this example, depending on the width and center frequency used by node G for link (G_1, A_1), A may not have a possible width combination that ensures a minimum channel for its remaining interface (numbered 2, link to C). Thus, node G while coloring link (G_1, A_1) should make sure that A has at least a block for its remaining uncolored interfaces. A valid channel allocation then is a channel allocation that is not only feasible from a node's perspective but also is consistent with channel allocations at other nodes.

Based on the above, we seek a channel allocation that results in a valid channel width combination at every node in the network. Before going to the actual algorithm, we need to introduce two more concepts: *guard block assignments* (*gBAs*) and *guard spectrum allocation maps* (*gSAMs*). Recall from constraint **C2** in Section 3 that we require each link in the network to be assigned at least a

---

[6]Relative preference among channel width combinations at a node can be determined, for example, by computing a function (e.g., product, sum) of *link utilizations* with each combination and then ranking the combinations based on their function values. Each link's utilization is computed by taking the ratio of its load to capacity, the latter based on channel width corresponding to the link in the chosen width combination. We use product of link utilizations as the default method in this paper.

[7]Henceforth we use the terms of coloring and channel allocation interchangeably.

minimum channel (block). We ensure that this constraint is met at all times (i.e., prior to, during and after completion of channel allocation algorithm execution) via $gBAs$ and $gSAMs$. A valid block is identified for each link at the initialization stage of the algorithm as described below in step 1. The block so identified for a link is referred to as the gBA for that link. The $gSAM$ of a node is a collective representation of $gBAs$ of all its incident links. Like $SAMs$, $gSAMs$ are also bit-vectors. After every link gets a gBA, it is straightforward to determine node $gSAMs$. For example, consider the network shown in Fig. 3 and suppose that the total amount of available spectrum is 40MHz with the block identifiers as in Fig. 2. A possible (though not optimal) gBA assignment for links G-A, A-C, C-B and B-G is blocks 1, 2, 3, 2 respectively. For that assignment, $gSAM_A$ is $< 1, 1, 0, 0, 0, 0, 0, 0 >$, $gSAM_B$ is $< 0, 1, 1, 0, 0, 0, 0, 0 >$ and so forth. The gBA assignments can be seen as proactively *reserving* a minimum sized channel for each link that results in a valid channel allocation to start with. Later on when the algorithm finds a larger width channel that is commensurate with the load on a link and does not compromise validity of the channel allocation, the gBA for that link is *released* in exchange for a valid block assignment corresponding to the larger width channel; $gSAMs$ and $SAMs$ of the end nodes of that link are also accordingly updated then.

Our channel width assignment algorithm consists of the following sequence of steps:

1. *Initialize $gBAs$ and $gSAMs$:* We do this by applying an edge coloring heuristic on the network, viewing each individual block in the given spectrum as a potential color. Even if this coloring is done greedily we are guaranteed to have a proper edge coloring given our assumption in Section 3 about the relationship between maximum degree of the network ($\Delta(\mathcal{T})$) and the total number of spectrum blocks (S). See [31]. Our assumption and choice of the heuristic are driven by the fact that greedy edge coloring can be easily implemented in a distributed manner [32].

2. *Assign node and link priorities based on traffic load:* The rest of the algorithm is also greedy, driven by priorities assigned to nodes and links. Node priorities determine the order in which nodes allocate channels to their interfaces (incident
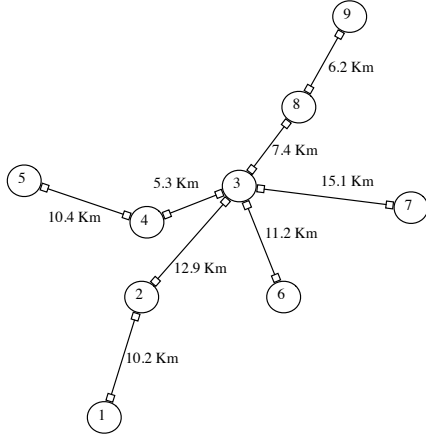
Figure 6: Aravind telemedicine network topology [2].

links) — higher a node's priority sooner its turn for channel allocation. Specifically, the priority of a node n, P_n is determined as follows:

$$P_n = \sum_{m, \exists (n_p, m_q)} f_{n_p, m_q} \qquad (1)$$

The above equation favors nodes with larger traffic volume when assigning priorities since the goal of the algorithm is to adapt based on traffic demands. Priority of a link (n, m) is set to the average of the priorities of end nodes, i.e., $P_{n_p, m_q} = avg\{P_n, P_m\}$.

3. *Steps taken when a node x is the next highest priority node to be processed:*

(a) Find the *list* of feasible width combinations at node x taking into account priorities of incident links at x and order the combinations based on their relative preference as described earlier in this section (see Fig. 4 and corresponding text).

(b) Prune the list from the previous step (3.a) to retain only potentially valid width combinations. Depending on the priority of node x relative to its neighbors, some of its incident links may already be colored by higher priority neighbors. In such a case, some of the combinations from (3.a) may not be valid causing their removal from the list. Fig. 5 illustrates this step.

(c) Find a valid channel width combination from the remaining list from step (3.b) considering combinations in the order of preference and stopping when a valid combination is found. First check for validity of a combination by verifying if the widths in the combination for uncolored incident links (x, y) at node x can be satisfied based on the current SAMs of x and all such neighbors y while not violating constraints **C1** and **C3**. If the first check is successful then the combination is checked for violation of gBAs for uncolored incident links (y, z) z ≠ x of neighbors y. If both these checks are successful, then search for a valid combination is successful. At that point, uncolored links of x are colored based on the combination found and block assign-

ments for those links and SAMs of x and affected neighbors y are updated. Moreover, gBAs for the newly colored links of x are released and gSAMs of end nodes accordingly updated. Note that this step will result in a valid combination being found because the combination corresponding to gBA assignment for the uncolored links of x prior to this step is always among the combinations searched during this step. Before finishing this sub-step, node x tries to move the existing gBAs to increase flexibility of remaining channel width assignments.

At the end of the execution of the above algorithm, we are guaranteed a valid channel allocation because validity of the channel allocation is maintained at each step of the algorithm. The fact that the algorithm in fact terminates and runs in polynomial time is also evident from the above description. At termination, gBAs for all links are released and gSAMs of all nodes become null vectors. Theoretically characterizing the approximation guarantee of this algorithm is an issue for future work. We discuss the practical aspects and distributed operation later in Section 6.

5. EVALUATION

In this section, we evaluate the effectiveness of our channel width assignment algorithm described in the previous section using simulations to understand its ability to adapt to spatio-temporal variations in traffic demands. Our goal is to study the benefit of adapting channel width across a diverse set of scenarios in terms of effective and fair allocation of the limited spectrum resource. Since we are not aware of channel width assignment algorithms for long-distance mesh networks, we conduct this study in comparison with a variant of [14] that does (the more common) undirected edge coloring. Within this benchmark, we consider several alternatives each based on a different fixed size channel width, starting from 5MHz (minimum sized channels in current 802.11 systems). We focus on rural wireless access network scenarios because they are a compelling real-world use case for long-distance 802.11 mesh networks.

We use the QualNet simulator that has a built-in detailed model for standard 802.11 CSMA/CA MAC protcol. We have added the variable channel width functionality to QualNet and validated it against the results reported in [6]. We use the Traffic-Gen application in the simulator for flexible realization of different traffic patterns with variable session durations and traffic loads. We use 1KB packets throughout. In all our experiments, we set the total available spectrum to 100MHz to match with the commonly used 5.8GHz band. For the bit-rate (modulation and coding scheme), we use 6Mbps unless mentioned otherwise. For the adaptive channel width case, the channel width assignment algorithm is executed periodically. Unless otherwise specified, we use a simulation length of 25 minutes with traffic flows starting after one minute.

For our evaluations, we use three real long-distance wireless network topologies:

1. Aravind telemedicine network in Southern India [2] consisting of 9 backhaul wireless nodes. Fig. 6 shows the topology of this network — node 3 is the hospital situated in a town and nodes 1, 5, 6, 7 and 9 are remote vision centers.

2. Connected Communities (ConCom) network [33] is a relatively large broadband wireless access network covering the Western Isles of Scotland with a population around 26,000 spread across 11 islands and span of over 200Km. This network consists of 34 backhaul sites interconnected by point-
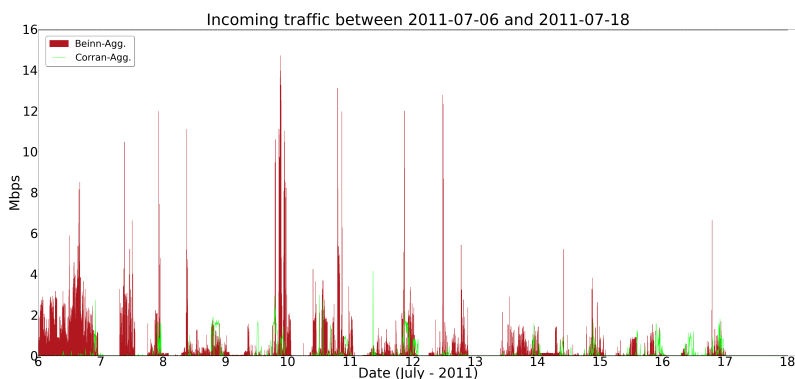
Figure 7: The trace of aggregated download traffic seen at two most used backhaul wireless sites in the Tegola network [13] over a two week period.
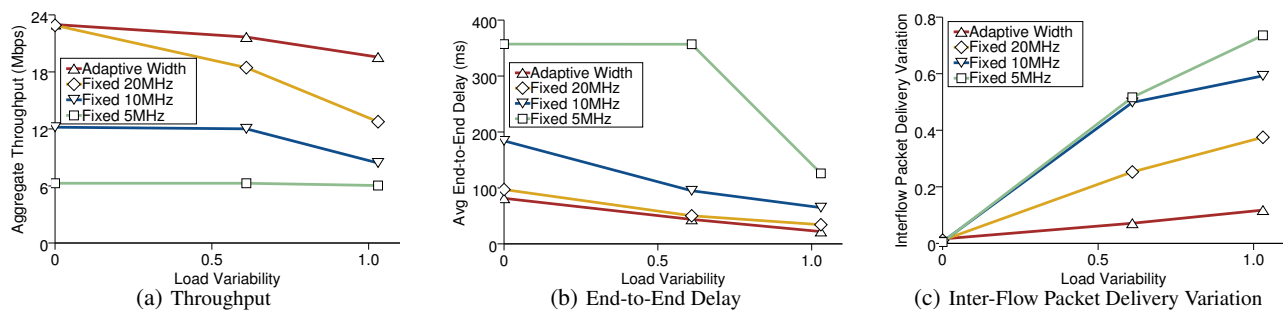

| (a) Throughput | (b) End-to-End Delay | (c) Inter-Flow Packet Delivery Variation |

Figure 8: Performance impact of spatial variation in traffic demands (variable load across flows) for the Aravind network topology.

to-point wireless links with widely different link lengths. It provides connectivity to public buildings (e.g., schools, community centers) as well as residential users.

3. Tegola consisting of 5 backhaul wireless nodes is a network we have deployed in rural Scotland [13]. Though originally intended as a research testbed, it currently also serves as a community wireless network connecting real users to the Internet. We have used this topology for our preliminary evaluations, but the results reported in this paper consider the other two relatively larger network topologies.

Before going to our simulation experiments, we first consider *a key aspect that motivates the need for adaptive spectrum management mechanisms, i.e., the spatio-temporal variability in traffic seen by different backhaul nodes in a long-distance wireless mesh network*. We illustrate the existence of such variability in real deployments in Fig. 7 with a 2-week trace of traffic at two Tegola network backhaul sites Beinn (shown in red) and Corran (shown in green) with most number of associated users (16 and 6, respectively). This trace corresponds to aggregated download traffic seen at each site; upload traffic (not shown) also exhibits similar behavior though as expected lower in volume compared to download. While variability in traffic across these sites and over time is visually evident, we also quantified this diversity by computing the pearson's correlation coefficient between the traffic at the two sites. This comes out to be 0.12, which indicates little or no correlation. In the following, we examine the performance of our channel width adaptation algorithm against such spatio-temporal traffic variations.

5.1 Spatial Traffic Variation

In the first experiment, we study the impact of spatial variation in traffic load. For this we consider the Aravind topology with 5 traffic flows from each of the edge nodes (1, 5, 6, 7, 9) to the middle node (3). We uniformly vary the traffic load across the flows while keeping the mean traffic load constant around 5Mbps. Fig. 8 shows the results for aggregate throughput, average end-to-end delay and packet delivery variation across flows as a function of load variability. Load variability is essentially the coefficient of variation[8] calculated using loads of individual flows. Packet delivery variation across flows shown in Fig. 8(c) is also a coefficient of variation but calculated using the packet delivery ratios of individual flows. For the Aravind network topology (as well as the ConCom topology), all links cannot be assigned 40MHz channels given the limited total amount of spectrum (100MHz), so only 20MHz, 10MHz and 5MHz fixed width allocations are shown as alternatives to adaptive channel width. As expected, the opportunity for adapting channel width is marginal when the traffic is uniform, while significant gains are achieved as load becomes more variable, resulting in throughput improvement around 53%. Improvements in variation of packet delivery across flows (Fig. 8(c)) at high load variability are even more remarkable as all the fixed width cases fail to support flows with high load. The big drop in delay for 5MHz fixed width case (Fig. 8(b)) can be explained by the fact that most of the high delay packets corresponding to high load flows are dropped. This is also reflected in the large variation of packet delivery across flows for this width in Fig.8(c).

[8]The ratio of standard deviation to the mean.
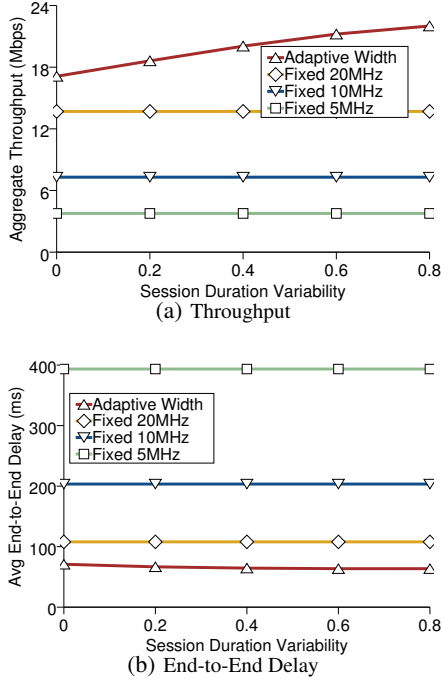
Figure 9: Performance impact of temporal variation in traffic demands (variable session durations across flows) for the Aravind network topology.

5.2 Temporal Traffic Variation

We now consider the impact of temporal variation in traffic demands, again using the Aravind network topology. For this experiment, we consider three flows (5-3, 6-3 and 7-3), each having the same load around 10Mbps but varying in session durations. Fig. 9 shows the throughput and delay results as a function of session duration variability, computed as the coefficient of variation using session durations of individual flows. Mean session duration was kept constant at 25 minutes for all the data points. Like in the previous spatial variation experiment, we find that adaptive width offers greater gains as temporal variability increases (up to 45% improvement in throughput) by adaptively reallocating spectrum as flows come and go.

5.3 Larger Network Scenario

We now examine the benefit of channel width adaptation in a larger network using the ConCom topology. For the traffic pattern, we consider a common use-case for the ConCom network, i.e., tele-commuting/tele-education. Sites connecting office buildings and schools act as traffic sources (8 in number) while traffic destinations for the flows are randomly distributed from among the remaining sites. We keep the load of each flow constant at around 10Mbps and increase the number of flows. Results are shown in Fig. 10. We observe that increasing the number of flows has the effect of making the traffic pattern more uniform, limiting the benefit of adaptive width. When the traffic pattern is non-uniform and less constrained by the amount of available spectrum (left half of the figure), adaptive width results in throughput improvement over 70% compared to the best fixed width alternative. The drop in delay from midway (especially for fixed width cases) seen in Fig. 10(b) is a result of packets with large delays getting dropped as contention increases and queues build up.
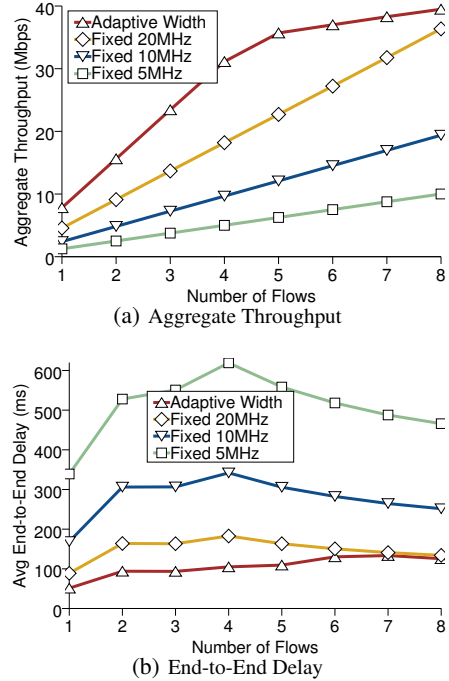
Figure 10: Performance impact from using adaptive channel width for the larger ConCom network with increasing number of flows.
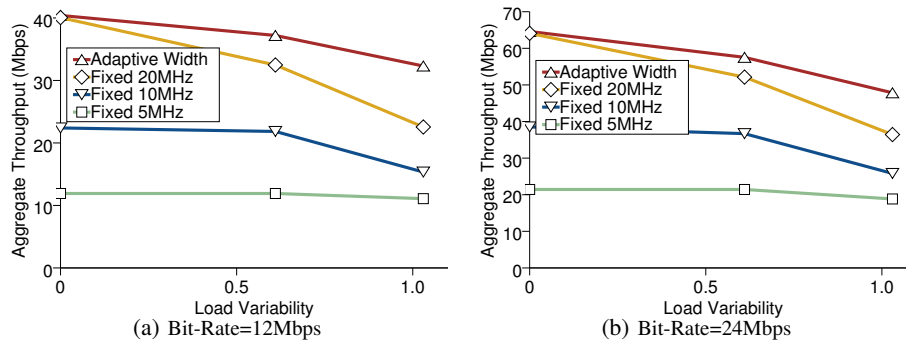
5.4 Effect of Bit-Rate

We also study the interaction between bit-rates and channel widths and their net effect on performance using the Aravind topology and considering two other bit-rates for the spatial variation experiment: 12Mbps and 24Mbps. The choice of these rates is based on the fact that 24Mbps is the maximum bit-rate that can be supported by link 3-7 for 40MHz. While the results shown in Fig. 11 qualitatively are similar to those in Fig. 8, latter corresponding to 6Mbps bit-rate, improvements from using adaptive channel width differ because the opportunity provided by use of increased widths somewhat reduces with increased bit-rates as also shown experimentally in Fig. 2 of [6].

6. DISCUSSION

For small to medium scale scenarios like the ones considered in this paper, channel width adaptation can be carried out in a centralized fashion at a gateway node. Gateway in such an implementation acts as a channel allocation server with each node periodically reporting measured link level traffic volume information to the server. Every adaptation interval, the channel allocation server uses that information and recomputes the new channel allocations for each link. The server then communicates them back to the network one node at a time, waiting for confirmation from the node that it completed channel reconfiguration locally through coordination with its neighboring nodes. This approach to implementing traffic-aware channel allocation is practical as demonstrated earlier for the omnidirectional mesh scenario by Ramachandran et al. [18].

Concerning the length of the adaptation interval itself, it depends on the traffic dynamics as well as network overhead for channel reallocation. Due to the aggregation of traffic from individual users at the backhaul nodes, the variability of traffic seen at a backhaul node over time is slower in the order of several minutes; this ob-

Figure 11: **Effect of inter-dependence between channel width and bit-rate on performance for the case of spatial traffic variation with Aravind network topology.**
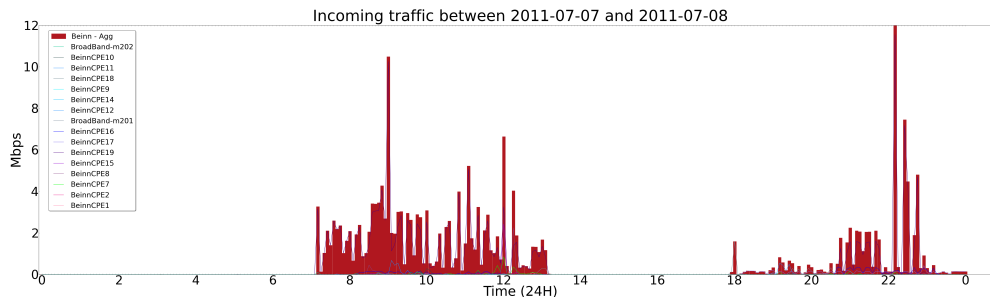
Figure 12: **The trace of aggregated and individual user download traffic seen at Beinn backhaul site in the Tegola network [13] on a typical week day. Aggregated traffic is shown in red.**

servation is also confirmed by our traffic traces from the Tegola network (Fig. 12). Also in rural wireless networks such as Tegola, aggregate traffic patterns are reasonably predictable with highs and lows around the same time each day. This could be exploited to schedule global channel width adaptations a priori.

For larger networks like [5], distributed implementation is required for scalability reasons. Here we outline the distributed implementation of our proposed channel width assignment algorithm below. Every instantiation of the algorithm consists of three phases corresponding to the three high-level steps described in Section 4: (1) initial distributed edge coloring for guard block assignments; (2) localized node priority assignment based on locally exchanging load information and using node IDs for breaking ties; (3) committing a node's coloring decision after finding a valid width combination — this phase requires the use of a distributed mutual exclusion mechanism along the lines of [19] to limit the number of nodes in a local neighborhood that can concurrently update their channel allocation. We leave the detailed specification of this distributed algorithm and its prototype implementation for future work.

The channel width assignment algorithm can be extended to consider inter-channel separation as a means to minimise adjacent channel interference as follows: We define the minimum required separation distance between any two incident channels ch1 and ch2 as a number of 5MHz blocks (i.e., $(f_{ch2} - w_{ch2}/2) - (f_{ch1} + w_{ch1}/2) \geq 5 * d$, where $f_{ch2} > f_{ch1}$ and d is the separation distance). We refer to these blocks as the *padding blocks*. Each link is then allocated the required channel for transmission as well as the *padding blocks*, which are chosen to the right of the *transmission channel* (except when the transmission channel lies at the upper end of the available spectrum, in which case no padding channel is nec-

essary). Although this requires a larger number of total available blocks to ensure valid channel allocation, one block channel separation between co-incident links may be sufficient. This, however, should be experimentally validated by considering other factors, such as antenna separation.

Finally, the effect of frequency-dependant attenuation, which increases as the carrier frequency increases can be considered in the proposed algorithm as follows. As attenuation increases, the signal-to-noise plus interference ratio, which determines whether a packet is correctly decoded at the receiver, decreases. This causes links operating at higher frequencies to use lower data rates, which restricts the available capacity. To capture this, we consider the effect of frequency on the raw capacity in our algorithm by measuring the effective data rate at each width using the lowest frequency in the available spectrum and use a coefficient to scale down capacity as the frequency increases. Moreover, the feasible width combinations at a node are ordered based on their relative preference assuming the center frequency of each width is the lowest possible frequency (step (3.a) of our algorithm in Section 4). Then, as each combination is checked for its validity (step (3.c)), its preference is scaled down based on the part of the spectrum each width actually occupies. At the end, a set of valid width combinations is formed among which the node chooses the one with the highest preference.

7. CONCLUSIONS

In this paper, we have studied the traffic-aware channel allocation in long-distance 802.11 mesh networks. We leverage the flexibility of using variable channel widths to adapt the channel allocation in response to spatio-temporal variations in traffic de-

mands. We show that the traffic-aware channel width assignment problem is NP-complete by establishing a relationship with the well known edge coloring problem. Our proposed polynomial time greedy heuristic algorithm results in a valid channel allocation for every node. Our simulation based evaluation of the algorithm using real network topologies shows that it substantially improves network performance (e.g., up to around 70% throughput improvement) relative to the existing fixed width allocation approach. As part of future work, we plan to develop a prototype implementation of the algorithm and conduct experimental evaluation over our Tegola network. Theoretical performance characterization of the proposed algorithm is another issue for future work. We also intend to come up with a detailed design of distributed channel width assignment protocol for long-distance 802.11 mesh networks.

8. REFERENCES

[1] Wireless Networking in the Developing World. http://wndw.net/index.html, Dec 2007. Second Edition.

[2] S. Surana, R. Patra, S. Nedevschi, M. Ramos, L. Subramanian, Y. Ben-david, and E. Brewer. Beyond Pilots: Keeping Rural Wireless Networks Alive. In *Proc. USENIX NSDI*, 2008.

[3] W. Tu, C. J. Sreenan, C. T. Chou, A. Misra, and S. Jha. Resource-Aware Video Multicasting via Access Gateways in Wireless Mesh Networks. In *Proc. IEEE ICNP*, 2008.

[4] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *Proc. ACM MobiCom*, 2005.

[5] NGI SpA. http://www.ngi.it/.

[6] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. In *Proc. ACM SIGCOMM*, 2008.

[7] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan. Load-Aware Spectrum Distribution in Wireless LANs. In *Proc. IEEE ICNP*, 2008.

[8] S. Rayanchu, V. Shrivastava, S. Banerjee, and R. Chandra. FLUID: Improving Throughputs in Enterprise Wireless LANs through Flexible Channelization. In *Proc. ACM MobiCom*, 2011.

[9] B. Raman. Channel Allocation in 802.11-based Mesh Networks. In *Proc. IEEE INFOCOM*, 2008.

[10] P. Dutta, S. Jaiswal, and R. Rastogi. Routing and Channel Allocation in Rural Wireless Mesh Networks. In *Proc. IEEE INFOCOM*, 2007.

[11] S. Nedevschi, R. K. Patra, S. Surana, S. Ratnasamy, L. Subramanian, and E. Brewer. An Adaptive, High Performance MAC for Long-Distance Multihop Wireless Networks. In *Proc. ACM MobiCom*, 2008.

[12] R. Patra, S. Nedevschi, S. Surana, a. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High PerformanceWiFi Based Long Distance Networks. In *Proc. USENIX NSDI*, 2007.

[13] G. Bernardi, P. Buneman, and M. K. Marina. Tegola Tiered Mesh Network Testbed in Rural Scotland. In *Proc. ACM MobiCom Workshop on Wireless Networks and Systems for Developing Regions (WiNS-DR'08)*, Sep 2008.

[14] P. Dutta, S. Jaiswal, D. Panigrahi, and R. Rastogi. A New Channel Assignment Mechanism for Rural Wireless Mesh Networks. In *Proc. IEEE INFOCOM Mini-Conference*, 2008.

[15] W. Yuan, W. Liu, and W. Cheng. Capacity Maximization for Variable-Width WLANs: A Game-Theoretic Approach. In *Proc. IEEE ICC*, 2010.

[16] R. Gummadi, R. Patra, H. Balakrishnan, and E. Brewer. Interference Avoidance and Control. In *Proc. ACM HotNets*, 2008.

[17] J. G. Andrews, A. Ghosh, and R. Muhamed. *Fundamentals of WiMAX*. Prentice Hall, 2007.

[18] K. Ramachandran, E. Belding, K. Almeroth, and M. Buddhikot. Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks. In *Proc. IEEE INFOCOM*, 2006.

[19] B. J. Ko, V. Misra, J. Padhye, and D. Rubenstein. Distributed Channel Assignment in Multi-Radio 802.11 Mesh Networks. In *Proc. IEEE WCNC*, 2007.

[20] S. Pediaditaki, P. Arrieta, and M. K. Marina. A Learning-based Approach for Distributed Multi-Radio Channel Allocation in Wireless Mesh Networks. In *Proc. IEEE ICNP*, 2009.

[21] M. F. Uddin, H. Alazemi, and C. Assi. Joint Routing, Scheduling and Variable-Width Channel allocation for Multi-hop WMNs. In *Proc. IEEE ICC*, 2010.

[22] L. Li, C. Zhang, and Y. Li. QoS-Aware On-Demand Channel Width Adaptation Protocols for Multi-Radio Ad-hoc Networks. In *Proc. IEEE WCNC*, 2009.

[23] F. Wu, N. Singh, N. Vaidya, and G. Chen. On Adaptive-Width Channel Allocation in Non-Cooperative, Multi-Radio Wireless Networks. In *Proc. IEEE INFOCOM*, 2011.

[24] S. M. Das, H. Pucha, D. Koutsonikolas, Y. C. Hu, and D. Peroulis. DMesh: Incorporating Practical Directional Antennas in Multichannel Wireless Mesh Networks. In *IEEE JSAC*, 2006.

[25] V. Angelakis, N. Kossifidis, S. Papadakis, V. Siris, and A. Traganitis. The Effect of Using Directional Antennas on Adjacent Channel Interference in 802.11a: Modeling and Experience With an Outdoors Testbed. In *Proc. IEEE WiNMee Workshop*, 2008.

[26] D. Tyrode. On the Predictability of Traffic and Energy Resource Availability for Self-Powered Wireless Network Routers. Master's thesis, The University of Edinburgh, Aug 2011.

[27] A. Sheth, S. Nedevschi, R. Patra, S. Surana, and E. Brewer. Packet Loss Characterization in WiFi-based Long Distance Networks. In *Proc. IEEE INFOCOM*, 2007.

[28] D. Gokhale, S. Sen, K. Chebrolu, and B. Raman. On the Feasibility of the Link Abstraction in (Rural) Mesh Networks. In *Proc. IEEE INFOCOM*, 2008.

[29] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[30] V. Raghunathan and P. R. Kumar. Wardrop Routing in Wireless Networks. In *IEEE Transactions on Mobile Computing*, 2009.

[31] D. B. West. *Introduction to Graph Theory*, chapter 7. Prentice Hall, 2000.

[32] M. Kodialam and T. Nandagopal. Characterizing Achievable Rates in Multi-hop Wireless Networks: The Joint Routing and Scheduling Problem. In *Proc. ACM MobiCom*, 2003.

[33] Connected Communities Network. http://www.connectedcommunities.co.uk/.