Traffic-Aware Channel Assignment for Multi-radio Wireless Networks*

Ryan E. Irwin^{1,2}, Allen B. MacKenzie¹, and Luiz A. DaSilva^{1,3}

¹ ECE Department, Virginia Tech, Blacksburg, Virginia U.S.A.

² Raytheon BBN Technologies, Cambridge, Massachusetts, U.S.A.

³ CTVR, Trinity College, Dublin, Ireland

{reirwin,mackenab,ldasilva}@vt.edu

Abstract. This paper studies channel assignment in multi-hop wireless networks in which nodes are equipped with multiple radios, each of which can be assigned to a channel. We argue for an approach that first assigns channels independently of traffic, to achieve basic connectivity and support light loads such as control traffic, and then dynamically assigns channels to the remaining radios in response to traffic demand. The objective is to balance the need for a stable baseline topology and the desire to maximize throughput by dynamically adapting the topology to current network conditions. We call this a traffic-aware (TA) approach, in contrast to both traffic-independent (TI) and traffic-driven (TD) channel assignment and topology control schemes found in the literature. We formulate the problem as a two-stage mixed integer linear program (MILP), and find that our approach supports good connectivity and data rates comparable to those achieved with a TD channel assignment, while achieving lower resource utilization than TI approaches. We also quantify the tradeoffs involved in the decision of what proportion of radios to enable during the traffic-independent stage and what proportion to enable dynamically in response to changing traffic demands.

1 Introduction

Although there has been tremendous success in the adoption of single-hop wireless networks (e.g. cellular networks and wireless local area networks (WLANs)), multi-hop wireless networks (e.g. ad hoc and mesh networks) struggle to gain widespread deployment. The first attempts at multi-hop wireless networking were naturally in a single-channel environment. However, spanning multiple hops on a single channel is difficult due to intra-flow contention and hidden terminals.

Using multiple orthogonal channels with multiple radios per node improves the viability of multi-hop wireless networks, but careful design of a channel

* This research was partially supported by a Bradley Fellowship from Virginia Tech's Bradley Department of Electrical and Computer Engineering and made possible by an endowment from the Harry Lynde Bradley Foundation. This material is also based upon works supported by the Science Foundation Ireland under Grant No. 10/IN.1/I3007.

assignment scheme is necessary. There are two critical design goals in channel assignment for multi-channel, multi-radio, multi-hop networks: (1) establishing and maintaining network connectivity, and (2) adapting the topology to meet the traffic demand. Since the goal of a network is to serve traffic, adapting the topology directly improves network utility (e.g. flow throughput). We argue that these two design goals be addressed together, since they both access the same pool of resources. An assignment scheme of one type of allocation limits the actions and effectiveness of the other.

In this paper, we explore the fundamental tradeoff in the balance of resources allocated traffic-independently and resources allocated in response to traffic demands. In the process, we propose a balanced scheme that we call a traffic-aware (TA) approach. The scheme allocates a minimum amount of resources traffic-independently. This supports the network's connectivity needs (goal 1), while allowing the network to adapt dynamically to changing traffic demands with the remaining resources (goal 2). The proposed traffic-aware (TA) allocation follows the ideas of cognitive networking [12] where the network senses its operating conditions and adapts its resource allocation accordingly.

The evaluation in this paper presents a broad comparison with two alternative types of resource allocation found in the literature. Some works propose a traffic-independent (TI) channel assignment, while others propose a traffic-driven (TD) channel assignment. In [7, 10, 11] the main objective is to minimize interference subject to obtaining a prescribed level of network connectivity. The schemes strive to assign channels to all radios of nodes in the network. Similarly, in [2, 6], all radios in the network are assigned channels with the objective of maximizing connectivity subject to a limited level of interference. These approaches all address the problem of achieving network connectivity independently of traffic conditions. (See [5] for a comparison of aforementioned schemes.) By contrast, there are other proposals in the research literature that address the resource allocation problem from a traffic-driven perspective, with the assumption of a zero-cost control plane enabling network connectivity. These include [1, 3, 4, 8, 9], which propose purely TD allocation schemes.

The remainder of this paper is organized as follows. Section 2 provides the two-stage mixed-integer linear program (MILP) formulations used to evaluate the soundness of this approach and explore the fundamental tradeoffs involved. Section 3 outlines our numerical procedure and analyzes the results. Section 4 concludes the paper.

2 Problem Formulation

We formulate the problem as a two-stage MILP of channel assignment and flow routing. Figure 1 shows the basic idea of the two-stage MILP. The first (TI) stage assigns channels, inducing a topology that does not change based on traffic conditions. The second (TD) stage assigns channels to any unused radios from the first stage, based on traffic conditions, with the goal of maximizing throughput by modifying the topology to be more suited to the current traffic

Fig. 1. Two-Stage MILP Approach: colors represent channels, solid and dashed lines represent TI and TD links respectively

demand. Under changing traffic conditions, the topology can be re-optimized to meet the current traffic demands. In this section, we define the two-stage MILP that models the TA approach that we propose. Subsequently, we define another two-stage MILP that is used for comparison in our evaluation of various types of resource allocation schemes.

2.1 Two-Stage, Traffic-Aware (TA) MILP

We propose, as our main contribution, an approach that allocates minimal resources in the first (TI) stage, with the goal of using as few radios as possible to create a connected topology, leaving as many unassigned radios as possible to optimize the topology based on traffic characteristics in the second (TD) stage. In the following subsections, we provide the details of the two-stage formulation.

Resource-minimized, TI Stage: Problem $\mathcal{RM}$. Starting with physical layer constraints, we assume the available spectrum is divided into orthogonal channels contained in set $\mathcal{C}$, and each radio occupies a single channel at any given time. Although radios can occupy any channel, we assume the cost of channel switching is too high to incur on a per-packet basis, so nodes cannot receive on one channel and transmit on another with the same radio. Nodes are contained in set $\mathcal{V}$. The objective function is

$$\min \sum_{i \in \mathcal{V}} |\mathcal{C}_i|, \quad (1)$$

where $\mathcal{C}_i$ is the set of channels tuned to by node i with its radios. Each node is equipped with K radios, so each node can occupy up to K channels. That is,

$$|\mathcal{C}_i| \leq K \quad (\forall i \in \mathcal{V}). \quad (2)$$

Also, we define $\mathcal{C}_{ij} = \mathcal{C}_i \cap \mathcal{C}_j$ as the set of channels nodes i and j have in common.

The communication and interference model we adopt for this problem is the double-disk model similar to [2, 5, 7, 10, 11]. There are two disks centered at each node. The inner disk, the communication range disk, has a radius of r_{comm} and the outer disk, the interference disk, has a radius of r_{int} ($r_{comm} < r_{int}$). A node i can communicate with any node j within its communication range if $\mathcal{C}_{ij} \neq \emptyset$. If node j is within node i 's interference disk and $\mathcal{C}_{ij} \neq \emptyset$, nodes i and j interfere with one another. By definition if two nodes can communicate with each other, they can interfere with each other as well.

We represent the double-disk relationships with sets $\mathcal{CR}_i$ and $\mathcal{IR}_i$ as the set of nodes within communication range and interference range respectively. We also define binary matrices as follows. The interference range matrix is IR , where $IR[i][j] = 1$ if node $i \in \mathcal{IR}_j$ and 0 otherwise. Similarly, we define the communication range matrix $CR[i][j] = 1$ if node $i \in \mathcal{CR}_j$ and j are in communication range and 0 otherwise. If $CR[i][j] = 1$, then $IR[i][j] = 1$ by definition.

We define communication graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where $\mathcal{E}$ is the edge set. An edge e_{ij} exists in $\mathcal{E}$ if $\mathcal{C}_{ij} \neq \emptyset$ and $CR[i][j] = 1$. We define set $\mathcal{P}_{ij}$ as the set of all paths between node i and j , where a path p_{ij} is a list $i, e_{ik}, k, e_{kl}, l, \dots, m, e_{mj}, j$ of nodes and edges where no node or edge repeats. Since we are trying to achieve basic network connectivity using minimal resources, we stop allocating resources once there is a path between all node pairs in the graph. That is,

$$|\mathcal{P}_{ij}| \geq 1 \quad (\forall i \in \mathcal{V}, \forall j \in \mathcal{V}), \quad (3)$$

For the first (TI) stage of the MILP we seek a channel assignment that results in a limited number of interferers (β) for each node. That is,

$$\sum_{j \in \mathcal{V}, j \neq i} (IR[i][j] - CR[i][j]) \cdot |\mathcal{C}_{ij}| \leq \beta \quad (\forall i \in \mathcal{V}). \quad (4)$$

Note that under certain parameters (i.e., a low value of β , a particularly awkward node placement, few radios, and/or few channels) this constraint can cause the problem to be infeasible. Although we set $\beta = 0$ in our numerical analysis, we evaluate scenarios where $|\mathcal{C}| = 10$ and $K \geq 2$. This reduces the probability of an infeasible problem, but if this constraint drives the problem to be infeasible, this constraint could be relaxed to allow *some* interferers by setting $\beta > 0$. In the scenarios we evaluated, setting $\beta = 0$ did not make the problem infeasible.

Traffic-driven (TD) Stage. After the TI stage completes, there is a traffic-independent, resource-minimized, multi-channel connected topology. Then, in the TD stage, channels are assigned to any radios not assigned in the TI stage, adapting the topology to maximize flow rate r . We denote the set of flows as $\mathcal{F}$. Each flow f has a weighting factor of w_f , which is used to scale the rate of each flow with respect to r . The objective function reads

$$\max \sum_{f \in \mathcal{F}} w_f \cdot r. \quad (5)$$

As in the first stage, we allow each node to occupy up to K (the number of radios per node) channels. We define $l_{ij}(c)$ as the physical layer transmission rate from node i to node j on channel c . We assume each radio has a maximum rate γ for the sum total of transmission and reception rate per channel. This constraint implies nodes can receive and transmit traffic on each channel/radio through some sort of time multiplexing. We also limit the aggregate physical layer link transmission rate on each channel within every interference disk to be γ . This implies that channel sharing (through time multiplexing) is possible among nodes co-located in an interference disk. These constraints read as follows.

$$\sum_{j \in \mathcal{V}, j \neq i} l_{ij}(c) + \sum_{j \in \mathcal{V}, j \neq i} l_{ji}(c) \leq \gamma \quad (\forall i \in \mathcal{V}, \forall c \in \mathcal{C}_i) \quad (6)$$

$$\sum_{j \in \mathcal{V}} \sum_{k \in \mathcal{V}, k \neq j} l_{jk}(c) \cdot IR[i][j] \leq \gamma \quad (\forall i \in \mathcal{V}, \forall c \in \mathcal{C}_i) \quad (7)$$

Each interference disk region is centered at node i , so the aggregate rate of all transmitters, j , within range of i is constrained. Note that these constraints do not enforce a strict scheduling; doing so would make the problem much more difficult to solve.

We constrain the aggregate communication rate of flows between nodes as:

$$\sum_{f \in \mathcal{F}} t_{ij}(f) \leq \sum_{c \in \mathcal{C}_{ij}} l_{ij}(c) \cdot CR[i][j] \quad (\forall i \in \mathcal{V}, \forall j \in \mathcal{V}, j \neq i),$$

where $t_{ij}(f)$ is the data transfer rate for flow f from node i to node j , irrespective of the set of channels involved. The source and destination nodes for flow $f \in \mathcal{F}$ are denoted as $s(f)$ and $d(f)$ respectively. Flow conservation at each intermediate node $i \in \mathcal{V} \setminus \{s(f), d(f)\}$, $f \in \mathcal{F}$ is stated as,

$$\sum_{j \in \mathcal{V}, j \neq i} t_{ij}(f) = \sum_{k \in \mathcal{V}, k \neq i} t_{ki}(f) \quad (\forall f \in \mathcal{F}, \forall i \in \mathcal{V}, i \neq s(f), d(f)).$$

Flow conservation out of each source node i is stated as,

$$\sum_{j \in \mathcal{V}, j \neq i} t_{ij}(f) = r \quad (\forall f \in \mathcal{F}, i = s(f)). \quad (8)$$

Since there is flow conservation at source and intermediate nodes, there is flow conservation at the destination node. Lastly, we have non-negativity constraints for flow and link rates,

$$t_{ij}(f) \geq 0 \quad (\forall i \in \mathcal{V}, \forall j \in \mathcal{V}, \forall f \in \mathcal{F}), \quad (9)$$

$$l_{ij}(c) \geq 0 \quad (\forall i \in \mathcal{V}, \forall j \in \mathcal{V}, \forall c \in \mathcal{C}). \quad (10)$$

2.2 Parameterized Formulation

In this subsection, we provide a similar formulation following the workflow outlined in Figure 1. The motivation for this formulation is to evaluate the fundamental tradeoff between the TI connectivity and flow rate. In this formulation,

we exchange the TI stage ($\mathcal{RM}$) for a connectivity-maximized TI approach, which we denote as problem $\mathcal{CM}(\alpha)$, where α is the proportion of network radios allowed to be assigned ($0 \leq \alpha \leq 1$).

As compared to $\mathcal{RM}$, $\mathcal{CM}(\alpha)$ maximizes network connectivity using a specified proportion of the radios in the network, α , and is the dual problem to $\mathcal{RM}$.¹ The second (TD) stage is unchanged from subsection 2.1. In this subsection, we define $\mathcal{CM}(\alpha)$ and discuss how to vary α to effectively encompass the existing research literature.

Connectivity-maximized (CM), TI Stage. The motivation for this formulation is to evaluate the fundamental tradeoff between the TI network connectivity and maximum flow rate. The objective is to maximize network connectivity using no more than $\alpha \cdot |V| \cdot K$ radios.

The usual graph-theoretic metric for graph connectivity is k -node-connectivity, where k is the minimum number of node-disjoint paths over all node pairs. We denote the number of node-disjoint paths between nodes i and j as P_{ij}^{ND} , which equals the cardinality of the maximum-sized subset of $\mathcal{P}_{ij}$ whose elements, which are paths, have no nodes in common other than nodes i and j . The drawback of the k -node-connectivity metric is that it is dominated by the worst node-pairing. We suggest a more granular metric k' , where k' is k plus the proportion of all node pairs $\{(i, j) : i, j \in \mathcal{V}, i \neq j\}$ with more than k node-disjoint paths $P_{ij}^{ND} > k$. This is written as:

$$k' = \frac{1}{|V| \cdot (|V| - 1)} \cdot \sum_{i \in V} \sum_{j \in V, j \neq i} \min(P_{ij}^{ND}, k + 1). \quad (11)$$

Note that expression $\min(P_{ij}^{ND}, k + 1) \in \{k, k + 1\}$ so k' is in the interval $[k, k + 1]$.

The objective in $\mathcal{CM}(\alpha)$ is maximizing k' . Constraints from Equations (2) and (4) from $\mathcal{RM}$ are included in $\mathcal{CM}(\alpha)$. As before, we constrain that each node can occupy up to K channels, but the proportion of all radios' assigned channels must not exceed α . That is,

$$\sum_{i \in V} |C_i| \leq \alpha \cdot |V| \cdot K. \quad (12)$$

Implications of α . As we vary α , the proportion of resources assigned independently of traffic conditions, between 0 and 1, we note the particular significance of $\alpha = 0$ and $\alpha = 1$. When $\alpha = 1$ all resources are allocated in a strictly TI manner as is proposed in [2, 6, 7, 10, 11]. The TI stage maximizes TI connectivity, but the second (TD) stage reduces to a linear program (LP) of flow routing since no radio variables remain free to be assigned in response to traffic conditions.

¹ The objective switches from minimizing the sum of $|C_i|$ terms in $\mathcal{RM}$ to maximizing k' in $\mathcal{CM}(\alpha)$. The constraints of connectivity from $\mathcal{RM}$ appear in the objective of $\mathcal{CM}(\alpha)$, and the terms from the objective function in $\mathcal{RM}$ (C_i terms) appear in the constraints of $\mathcal{CM}(\alpha)$.

On the other extreme, when $\alpha = 0$ resource allocation is completely TD as is proposed in [1, 3, 4, 8, 9] and the TI stage is bypassed. The flow rate with $\alpha = 0$ represents the optimal flow rate since all resources are allocated in response to traffic conditions, so this extreme bounds all other solutions with $\alpha > 0$, including our TA scheme. This approach establishes network connectivity purely in reaction to the current traffic load and does not guarantee a connected network that can be used to support control traffic that is independent of current flows.

3 Numerical Investigation

We have used MATLAB and CPLEX to execute both stages of the two-stage MILP. As outlined in Figure 1, the output of the first (TI) stage of the MILP is input into the second (TD) stage. The second stage is repeated upon changes in the traffic demand.

First, we present an example network, showing the channel assignment response to a set of flows. This illustrates how the types of allocation differ from one another. Second, using multiple simulations, we show that the flow rate of the TA scheme tends to be significantly higher than the completely TI scheme and approaches the optimal throughput as provided by the completely TD scheme in 3- and 4-radio scenarios. Third, we illustrate the tradeoff by sweeping α in the interval $(0, 1)$ and show that the TA approach achieves an appropriately balanced allocation, providing enough connectivity while achieving sufficient flow rate as compared to the upper-bound.

3.1 Example Two-Stage Allocations

This example involves a uniformly random placement of 12 nodes with 2 radios per node and 5 channels. Range r_{comm} is set to 1.1 in a rectangular area of size 2 by $\frac{1}{2}$, and rate parameter $\gamma = 1$. Three source-destination node pairs are chosen at random according to a uniform distribution. Each flow has weight $w_f = 1, \forall f \in \mathcal{F}$.

Figure 2 shows the channel assignment of the three approaches. Each subfigure represents one of the approaches. Each line segment represents a communication link between two nodes, with each shade of grey representing a different channel. Flows are indicated by black circles with shaded centers, and each pair of circles of the same shade represents a source-destination pair. Links enabled through TI channel assignment are solid, and links that are enabled through TD channel assignment are dashed, and each node that assigns channels in response to traffic conditions is highlighted with a black triangle.

In the example illustrated in Figure 2, the TA scheme supports two thirds more flow rate of the TI scheme, despite having lower initial network connectivity. Also, the TA scheme achieves the same flow rate as the TD scheme, which is flow-rate optimal, while supporting a connected topology, which the TD approach does not.

Fig. 2. Assignments for a Single Set of Flows for each Approach

3.2 Results for TA, TD ($\alpha = 0$), and TI ($\alpha = 1$)

We run multiple simulations to characterize the flow rates using each approach. We use a rectangular area with edges of length 2 and 0.5, and we vary r_{comm} , effectively varying node density, to be 0.5, 0.8, and 1.1 with $r_{int} = 1.75 \cdot r_{comm}$. The variation in communication range causes the network diameter to vary from 2 to 6. Using a rectangle, as opposed to a square, allows us to see 6-hop network diameters without increasing the number of nodes too much.

We vary the number of radios per node (K) to be 2, 3, and 4 and the number of channels ($|\mathcal{C}|$) is 8. We use multiple lay-downs of nodes at each value of r_{comm} and K , with multiple flow configurations for each lay-down and 4 flows per configuration. We select source-destination node pairs randomly according to a uniform distribution. By varying the source-destination node pairings, we encompass many traffic scenarios. To avoid flow starvation, the weights of all flows are equal, enforcing flow fairness. Unbalanced flow weighting is left as future work. Flow rates are measured with respect to the radio capacity parameter $\gamma = 1$.

Figure 3a shows the flow rates while Figure 3b shows the radio usage pattern. Both figures show results for the scenarios where $r_{comm} = 0.8$. When $K = 3$ or $K = 4$, the flow rate supported by the TA scheme is almost identical to the flow rate of the TD scheme. In the 2-radio scenarios, the TA scheme achieves within 25% of the flow rate of the TD approach. This occurs because solving $\mathcal{RM}$ 1.2 radios per node are used on average to establish basic connectivity, leaving fewer radios available for a traffic-driven assignment.

The TA scheme achieves an average flow rate of 126%, 75%, and 80% more than that of the TI scheme in 2-, 3-, and 4- radio scenarios respectively. This significant gain shows how a TA assignment can outperform a TI scheme, supporting our assertion that a channel assignment should be dynamic, influenced by the traffic conditions. In the 2-radio scenario, fewer radios are conserved for a TD allocation (0.8 per node on average), but there is still a 126% increase in flow rate, suggesting that even modest TD resource allocation is beneficial.

Fig. 3. In Subfigure (b) schemes are numbered on the x-axis in each plot in order: (1) TI, (2) TA, and (3) TD. Term ‘Static’ refers to radios assigned traffic-independently. The term ‘Remain in-use’ refers to radios that remain on any channel when flow sources and destinations change. The terms ‘Newly-enabled’ and ‘Newly-disabled’ refer to radios turned on and off, respectively, when flow sources and destinations change.

The trends are similar for scenarios with $r_{comm} = 0.5$ and $r_{comm} = 1.1$ with the flow rates scaling based on the flow’s length in hop count. When $r_{comm} = 0.5$ networks support 25% less flow and when $r_{comm} = 1.1$ networks support 18% more flow. This is due to needing more resources to serve traffic over more hops.

In addition to achieving a higher flow rate, the TA scheme uses fewer radios, an energy-saving bonus. As compared to the TI scheme, it uses 15%, 36%, and 50% fewer radios when nodes are equipped with 2-, 3-, and 4- radios, respectively, on average. We track how the TA scheme and the TD scheme adapt their assignments to achieve their respective maximum flow rates. There are 40% percent fewer radios used in the TD scheme as compared to the TA scheme due to our connectivity requirements placed on the TA scheme. However, the TD scheme has 76% more radios that are newly-enabled and 79% more newly-disabled radios on average, which represents more fluctuations in the topology. In practice, such fluctuations induce control traffic at multiple layers, which takes away from network goodput and leaves the network vulnerable to control traffic storms.

As before, the results when $r_{comm} = 0.5$ and $r_{comm} = 1.1$ were similar to the scenario with $r_{comm} = 0.8$. The minor difference is that the TA scheme allocates more resources independently of traffic conditions when r_{comm} was lower and fewer resources when r_{comm} was higher. This occurs because more resources are required to support a connected topology when the radio range is lower.

3.3 Effects of Variable α

We have shown the benefits of the proposed TA allocation as compared to a completely TI approach ($\alpha = 1$) and a completely TD approach ($\alpha = 0$), but there is a tradeoff space where the proportion of resources allocated traffic-independently, α , can fall between 0 and 1. Note that $1 - \alpha$ is the proportion of resources available to be allocated in response to traffic demand. As α grows, then

fewer resources are available to adapt the network to changing traffic demand. At the other end of the spectrum, if α is too low, the network may not have sufficient baseline connectivity.

In this scenario, the parameters are the same as in the previous subsection. We have selected to display the scenario with $r_{comm} = 0.8$ because the trends for the scenarios with $r_{comm} = 0.5$ and $r_{comm} = 1.1$ were roughly equivalent.

Fig. 4. Relationships among α , k' , and r with $r_{comm} = 0.8$

Figure 4a shows the relationship between the proportion of resources allocated independently of traffic conditions (α) and the average k' -connectivity. We expect a monotonic increase in k' -connectivity as α increases, but we note the importance of how k' increases from $\alpha = 0$ to $\alpha = 1$. For each scenario (number of radios/node), there is a slow growth rate in k' initially with respect to α . Then, the slope of k' increases for values of α that correspond to a k' -connectivity in the range of $k' = (0.5, 2.0)$, indicating that there is a larger payoff in terms of connectivity for the resources allocated in this region. For larger values of α ($k' > 3$), the curve flattens, indicating that additional TI allocation of resources brings about diminishing payoff in terms of k' -connectivity.

Figure 4b shows α versus flow rate, r , plotted with 95% confidence intervals. We expect to see a decrease in rate as α increases, because as α increases the *flexibility* of the network's channel assignment decreases, meaning its dynamic response to traffic demand will have more limited effectiveness.

For low levels of α the flow rate does not decrease any as α grows, and all scenarios are roughly equivalent with respect to flow rate. The flow rate in the scenario with 2 radios per node drops off more quickly than in the scenarios with 3 and 4 radios per node. In each scenario, the rate of decrease in flow rate accelerates until $\alpha = 1$. The results show a flow rate plateau (especially for 3 and 4 radios per node), and to the right of the plateau the descent is steepest in scenarios with more radios per node. As α approaches 1, the gap between all flow

rates narrows. This was found to also be true at $\alpha = 0$, indicating that there is less flow rate benefit associated with increasing the number of radios at $\alpha = 0$ or $\alpha = 1$ as compared to other values of α . This further weakens the justifications for using strictly a TI or TD scheme in networks with multiple radios per node.

4 Conclusion

In this paper, we introduce a new approach to channel assignment in the context of multi-hop, multi-channel, multi-radio wireless networks and, more broadly, cognitive networks and dynamic spectrum access networks. We propose that networks employ a *traffic-aware* (TA) approach to channel assignment. The TA approach has two stages of channel assignment in which the first stage is a resource-minimized traffic-independent (TI) assignment and the second stage is a traffic-driven (TD) assignment. The TI assignment provides basic network connectivity with limited interference, enabling light traffic demands and basic network maintenance functions (e.g. serving control traffic). We propose allocating a minimum number of radios for the TI component to conserve as many resources as possible for the TD component, which is influenced by traffic demand. The conserved radios are dynamically assigned in response to traffic stimuli to maximize flow rate. As traffic conditions change, the TD assignments are re-optimized for current traffic conditions while the TI component remains unchanged.

Overall, results indicate that our TA scheme:

1. supports higher flow rate than a strictly TI scheme while using fewer radios,
2. approaches the flow rate of a strictly TD scheme in the 3- and 4-radio scenarios explored, and
3. has fewer fluctuations in channel allocations than a strictly TD allocation.

By finely varying the proportion of resources allocated independently of traffic conditions (α), we saw that k' -connectivity (a metric we define which is closely related to traditional k -connectivity) monotonically increases as a function α . Beyond a certain point, there are diminishing returns in terms of additional network connectivity gained through the allocation of additional radios for connectivity. Also, we see that flow rate r decreases as a function of α , and beyond a point the average flow rate diminishes more quickly. Both of these points suggest adopting a resource-minimized traffic-independent channel assignment in order to keep α as low as possible while still maintaining network connectivity.

As future work, we plan to analyze a set of scenarios where both TI and TD allocation decisions are performed locally at each node, building upon our previous work in [4] and [5]. In the process, we plan to develop a set of strategies nodes can follow to determine what constitutes a flow and how to determine if a local adaptation is beneficial for the network. We plan to compare the performance of the distributed approach with the (centralized) TA MILP in this paper, which represents a theoretical bound on the performance of any distributed approach. Also, we plan to investigate impacts of varying the traffic model to include more dynamic scenarios (e.g. vary the number of flows and flow lifetime).

References

- [1] Alicherry, M., Bhatia, R., Li, L.: Joint Channel Assignment and Routing for Throughput Optimization in Multi-Radio Wireless Mesh Networks. In: Proceedings of ACM International Conference on Mobile Computing and Networking (MOBICOM), p. 72. ACM (2005)
- [2] Chen, T., Zhang, H., Maggio, G., Chlamtac, I.: CogMesh: A Cluster-based Cognitive Radio Network. In: Proceedings of IEEE Dynamic Spectrum Access Networks (DySPAN), pp. 168–178 (2007)
- [3] Deaton, J., Ahmad, S., Shukla, U., Irwin, R.E., DaSilva, L.A., MacKenzie, A.B.: Evaluation of Dynamic Channel and Power Assignment for Cognitive Networks. *Wireless Personal Communications* 57, 5–18 (2011)
- [4] Irwin, R.E., DaSilva, L.A.: Channel Assignment Based on Routing Decisions (CARD): Traffic-Dependent Topology Control for Multi-Channel Networks. In: Proceedings of IEEE International Conference on Communications (ICC) CogNet Workshop (2009)
- [5] Irwin, R.E., MacKenzie, A.B., DaSilva, L.A.: Resource-Minimized Channel Assignment for Multi-transceiver Wireless Networks. In: Proceedings of IEEE Global Communications Conference, GLOBECOM (2011)
- [6] Jakllari, G., Ramanathan, S., Redi, J.K., Coffin, D.A., Tetteh, W.N., Burgess, J.R., Irwin, R.E.: Distributed Assignment of Frequency Channels to Transceivers over Dynamic Spectrum, US Patent App. 20,100/232,372 (September 16, 2010)
- [7] Marina, M.K., Das, S.R., Subramanian, A.P.: A Topology Control Approach for Utilizing Multiple Channels in Multi-radio Wireless Mesh Networks. *Computer Networks* 54(2), 241–256 (2010)
- [8] Neely, M., Ugaonkar, R.: Optimal Backpressure Routing for Wireless Networks with Multi-receiver Diversity. *Ad Hoc Networks* 7(5), 862–881 (2009)
- [9] Shi, Y., Hou, Y.: A Distributed Optimization Algorithm for Multi-hop Cognitive Radio Networks. In: Proceedings of IEEE International Conference on Computer Communications (INFOCOM), pp. 1292–1300. IEEE (2008)
- [10] Subramanian, A.P., Gupta, H., Das, S.R., Cao, J.: Minimum Interference Channel Assignment in Multiradio Wireless Mesh Networks. *IEEE Transactions on Mobile Computing* 7(12), 1459–1473 (2008)
- [11] Tang, J., Xue, G., Zhang, W.: Interference-aware Topology Control and QoS Routing in Multi-channel Wireless Mesh Networks. In: Proceedings of ACM International Symposium of Mobile Ad Hoc Networking and Computing (MOBIHOC), pp. 68–77. ACM (2005)
- [12] Thomas, R.W., Friend, D.H., DaSilva, L.A., MacKenzie, A.B.: Cognitive Networks: Adaptation and Learning to Achieve End-to-end Performance Objectives. *IEEE Communications Magazine* 44(12), 51–57 (2007)