Robust Design of Spectrum-Sharing Networks

Qingkai Liang, *Student Member, IEEE*, Hyang-Won Lee, *Member, IEEE*, and Eytan Modiano, *Fellow, IEEE*

Abstract—In spectrum-sharing networks, primary users have the right to preempt secondary users, which can significantly degrade the performance of underlying secondary users. In this paper, we use backup channels to provide reliability guarantees for secondary users. In particular, we study the optimal white channel assignment that minimizes the amount of recovery capacity (i.e., bandwidth of backup channels) needed to meet a given reliability guarantee, where both deterministic and probabilistic requirements are considered. This problem is shown to be coupled by two NP-hard objectives. We characterize the structure of the optimal assignment and develop bi-criteria approximation algorithms. Moreover, we investigate the scaling of the recovery capacity as the network size becomes large. It is shown that the recovery capacity is negligible as compared to the total traffic demands in a large-scale network.

Index Terms—Multi-hop spectrum-sharing networks, reliability guarantee, backup channel, channel assignment algorithms, minimum recovery capacity

1 Introduction

Sharing of radio spectrum is considered one of the most promising approaches to addressing spectrum scarcity in the face of exponentially growing traffic demands. Spectrum which is originally dedicated to a certain service can be shared based on access technologies such as cognitive radios [2] and geographic databases [3]. Typically, in spectrum-sharing networks, secondary users can access licensed channels (referred to as white channels) that are not used by primary users. While spectrum sharing enables efficient utilization of spectrum resources, secondary networks built upon white channels can suffer from severe performance degradation since secondary users must stop using a white channel whenever it is reclaimed by a primary user (this event is called *channel preemption*). Thus, it is necessary to provide protection for secondary users to guarantee their reliability against channel preemptions. In this paper, we study the problem of providing a reliability guarantee for multi-hop cognitive radio networks.

There have been numerous efforts towards achieving reliable communications for secondary users. One of the important issues in this context is that in the first place, it is desired for the transmitter and receiver to tune into the same channel, which preferably remains unoccupied by primary users for a long time. There are several works that exploit channel statistics to predict channel availability and design reliable MAC protocols for reducing the probability of being preempted. In [4], an design framework for MAC and sensing protocols is proposed that maximizes the performance of secondary network by learning channel statistics on-the-fly. In [5], a CSMA/CA-based data channel agreement and selection algorithm is presented in order to increase the rate of successful transmissions. Channel statistics are also exploited for application-aware channel selection. For example, in [6], real-time traffic is assigned channels with small variance of capacity, whereas elastic traffic is assigned channels with large capacity. Some works even apply machine learning techniques in order to decide on more stable channels [7], [8].

Another important issue is how the secondary network should recover from channel preemptions. A straightforward approach is to let disrupted links switch to another idle white channel on the fly [9], [10], [11]. This approach can, however, experience unpredictable delay until idle white channels become available. In contrast to the on-the-fly reconfiguration method, Yue et al. [12] propose to assign an extra white channel to each link in advance, in order to recover from any single channel preemption. In multi-hop networks, rerouting can be used to find a detour around interrupted links [13], [14] (see [15] for various rerouting methods and metrics capturing spectrum (and thus, link) availability). Some recent works [16], [17] combine channel switching and rerouting to recover secondary users' traffic. Moreover, there are some papers that study reliable channel assignments that maintain the network connectivity after any single channel preemption [18], [19], [20]. To better account for the performance of secondary network, the concept of spectrum leasing is proposed [21], where the primary network can lease part of its licensed spectrum and give priority to secondary network over the leased spectrum. Similarly, in [22], part of licensed channels are reserved for primary network to use in the event of channel failures as well as secondary network in the face of channel preemption due to primary network's activity. In these settings, the performance of both primary and secondary networks is dictated by the number (or portion) of leased channels, and hence, dynamic channel leasing/reservation schemes are employed so that the set of opened channels is changed adaptively to network status such as traffic intensity [21], [22].

- *Q. Liang and E. Modiano are with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: {qingkai, modiano}@mit.edu.*
- *H.-W. Lee is with the Department of Software, Konkuk University, Seoul 05029, Korea. E-mail: leehw@konkuk.ac.kr.*

Manuscript received 31 Jan. 2018; revised 9 July 2018; accepted 11 Aug. 2018. Date of publication 27 Aug. 2018; date of current version 28 June 2019.
(Corresponding author: Hyang-Won Lee.)
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TMC.2018.2867022

Although the above schemes enhance the reliability of secondary networks, most of them only provide "best-effort reliability". That is, there is no guarantee on, for example, the number of channel preemptions the secondary network can recover from, or the ability to fulfill a certain reliability requirement. It may be often the case that protection against single channel preemptions is not sufficient because simultaneous multiple channel preemptions can occur frequently due to the dynamic nature of white channels. It is also important to take into account the "schedulability" in the event of preemption. That is, the traffic demand may not be supported due to insufficient capacity, even if the connectivity is guaranteed.

In this paper, we allow secondary users to specify a reliability requirement and investigate how to adhere to such a requirement at the minimum cost. Our approach uses backup channels to recover from preemptions. These backup channels can be licensed channels leased temporarily at a cost [23], or currently unused white channels. Note that these backup channels do not necessarily stay idle when they are not used for recovery; the only requirement is that they should be available when needed for recovery (possibly at a cost).

Due to the scarcity and relative high costs of backup channels, it is desired to minimize the amount of *recovery capacity* (i.e., bandwidth of backup channels) that should be provisioned. Although many factors can affect the amount of required recovery capacity, we focus on the influence of *white channel assignment* to each link in the network. Specifically, we study the optimal white channel assignment that minimizes the recovery capacity required to meet a certain reliability requirement such that the network is able to recover secondary users' traffic from a given number of white channel preemptions. We consider both deterministic and probabilistic reliability requirements. Under the deterministic requirement, the network should be able to recover secondary users' traffic from a given number of white channel preemptions. Under the probabilistic requirement, secondary traffic should be recoverable from an arbitrary channel preemption scenario with a certain (high) probability.

Unfortunately, this problem is shown to be intractable and coupled by two NP-hard objectives. As a result, we conduct bi-criteria analysis and propose bi-criteria approximation algorithms for white channel assignment. For probabilistic recovery requirements, we also characterize the structure of the optimal channel assignment in different probability regimes of channel preemptions. Such an analysis leads to an efficient heuristic channel assignment scheme that can be used to meet any probabilistic recovery requirement. Our simulations validate the performance of the proposed algorithms.

It should be mentioned that although this work mainly focuses on the spare channel approach, our technical results can also be very useful for the schemes that use only white channels (such as the ones in [9], [10]). Note that even if white channels are used for recovery, the capacity of those idle white channels can possibly be insufficient to support disrupted traffic. Hence, it is important to construct a network in a way that the amount of disrupted traffic (or recovery capacity) is minimized, regardless of whether white channels or spare channels are used for recovery.

Another important contribution of this paper is the characterization of the scaling of the recovery capacity. It turns out that the required recovery capacity becomes negligible as compared to the total network traffic as the network becomes large. Our simulations show that under the proposed channel assignment schemes the required recovery capacity is usually less than 1 percent of the total traffic. Thus, it is possible to provision guaranteed reliability in a large-scale secondary network at minimum cost.

The remainder of this paper is organized as follows. We introduce the network model and describe the problem in Section 2. Next, we study the optimal white channel assignment under deterministic and probabilistic recovery requirements in Sections 3 and 4, respectively. Finally, simulation results are presented in Section 5 and conclusions are given in Section 6.

2 Model and Problem Description

2.1 Network Model

We consider a spectrum-sharing network where primary users own a set of licensed channels referred to as *white channels*. Any idle white channel can be accessed by secondary users, but it should be vacated if a primary user appears in that channel (referred to as *channel preemption*). When channel preemptions happen, secondary users switch to backup channels in order to resume communications. As mentioned in the introduction, the backup channels are assumed to be available when needed for recovery. Accordingly, there is no switching delay from white channel to backup channel. The *recovery capacity* refers to the bandwidth of backup channels we need to provision in order to meet a certain reliability requirement which will be specified in Section 3.

The secondary network is represented by an undirected graph $G = (V, E)$, where V is the set of secondary nodes and E is the set of links. There is a link between two secondary nodes if they can directly communicate with each other. We consider the one-hop interference model where adjacent links cannot be active on the same channel at the same time. Although such an interference model is restrictive, it serves as the foundation for understanding more complex interference models (e.g., see [24], [25]). Moreover, the one-hop interference model is an appropriate model for many practical wireless systems such as spread-spectrum systems, millimeter-wave networks [26], etc. Each link e is associated with a traffic demand r_e which is determined by some higher-layer policies (e.g., routing and flow control). We denote by W the set of white channels. Each white channel w can sustain a data rate up to $R_{w,e}$ over link e.

Now we describe the set of feasibility conditions on white channel assignment in order to sustain the given traffic demands. Let y be an $|E| \times |W|$ binary matrix whose element $y_e^w = 1$ if white channel w is assigned to link e. Note that if white channel w is assigned to link e, this link should be scheduled for at least $\frac{r_e}{R_{w,e}}$ fraction of time in order to meet the traffic demand r_e. Under the one-hop interference model, the set of links that can be activated simultaneously on the same channel form a *matching*, and interfering matchings can access the same white channel in a time-sharing manner. As a result, the set of feasible schedules can be represented by the *convex hull* of all matchings, i.e., the *matching polytope*. Based on Edmonds' matching polytope description [27], we can write the following feasibility conditions:

$$\sum_{e \in \delta(v)} \frac{r_e}{R_{w,e}} y_e^w \leq 1, \forall v \in V, w \in W \tag{1}$$

$$\sum_{e\in E(U)} \frac{r_e}{R_{w,e}} y_e^w \le \frac{|U|-1}{2}, \forall U \in \mathcal{V}, w \in W \quad (2)$$

$$\sum_{w\in W} y_e^w = 1, \forall e \in E$$
$$y_e^w \in \{0,1\}, \forall e \in E, w \in W. \quad (3)$$

In (1), we denote by $\delta(v)$ the set of links incident on node v. In (2), we define $\mathcal{V} = \{U \subseteq V : |U| \text{odd} \ge 3\}$ to be a collection of node sets with odd cardinality, and $E(U)$ is the set of links whose both ends are in U. For each white channel w, the corresponding constraints in (1) and (2) are Edmonds' matching polytope description over the set of links using that channel. Specifically, the constraints in (1) require the total schedule length of channel w not to exceed one; the constraints in (2) are called "odd-set constraints" and we refer readers to [28] or [27] for a detailed explanation. Overall, the constraints in (1) and (2) force all of the traffic demands to be schedulable under one-hop interference by using the given set of white channels. Hajek et al. [28] use a similar formulation to characterize schedulability in a single-channel case. Finally, the constraints in (3) force each link to be assigned exactly one white channel. A channel assignment y is said to be *feasible* if it satisfies all of the above constraints. The traffic demands $[r_e, \forall e \in E]$ are said to be *sustainable* under a channel assignment y if y is feasible.

2.2 Problem Description

Due to the scarcity and relatively high costs of backup channels, it is necessary to minimize the amount of recovery capacity (i.e., bandwidth of backup channels) needed to comply with a certain reliability requirement. In this paper, the secondary network is required to survive a given number of channel preemptions even in the worst case.

Given a recovery requirement, the amount of recovery capacity we need to provision depends on how much traffic is lost due to channel preemptions, which is largely determined by the assignment of white channels. For example, Fig. 1 illustrates two different channel assignments with 3 white channels. Each link has one-unit traffic demand, and we assume white channels have sufficiently large capacity such that any channel assignment is feasible (as long as each link is assigned exactly one white channel). Suppose we want to survive any single white channel preemption. In Fig. 1a, the preemption of channel 1 will cause the failures of two adjacent links, which requires two units of recovery capacity under one-hop interference. In contrast, the channel assignment in Fig. 1b only requires one unit of recovery capacity, since any links that can fail at the same time (i.e., due to the failure of a single white channel) can be activated simultaneously. *Our goal is to find a feasible white channel assignment that requires the minimum recovery capacity subject to a certain recovery requirement*.

3 Robust White Channel Assignment

In this section, we develop robust white channel assignment schemes that fulfill a given *deterministic* recovery requirement at minimum cost. Specifically, the network is required to survive any k white channel preemptions, i.e., the backup channels should be able to support the traffic demands on the links disrupted by any k white channel preemptions.

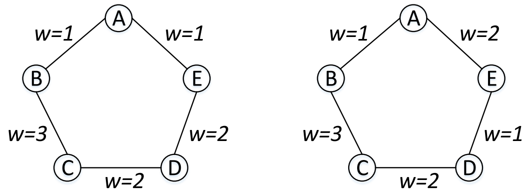

(a) Channel Assignment I (b) Channel Assignment II

Fig. 1. Two different assignments of white channels.

Hence, the goal is to find a feasible white channel assignment requiring the minimum recovery capacity to protect against any k channel preemptions. This problem is referred to as WhiteRec

$$\min_{C, y \text{ feasible}} \quad C$$
$$\text{s.t.} \sum_{w\in S} \sum_{e\in\delta(v)} \frac{r_e}{C} y_e^w \le 1, \forall v \in V,\ \forall S \in \mathcal{W}(k) \quad (4)$$

$$\sum_{w\in S} \sum_{e\in E(U)} \frac{r_e}{C} y_e^w \le \frac{|U|-1}{2}, \forall U \in \mathcal{V}, S \in \mathcal{W}(k), \quad (5)$$

where the meanings of $\delta(v)$, $E(U)$ and $\mathcal{V}$ are the same as in (1) and (2), and $\mathcal{W}(k) = \{S \subseteq W : |S| = k\}$ is a collection of channel sets with cardinality k. Similar to (1) and (2), the constraints in (4) and (5) correspond to Edmonds' matching polytope description, requiring that after any k channel preemptions the traffic demands on the disrupted links be schedulable by using a backup channel with capacity C. Hence, the optimal solution to WhiteRec is a feasible white channel assignment with the minimum recovery capacity. The constraints in (4) and (5) deterministically guarantee that the network survives any k channel preemptions. In Section 4, we study the probabilistic reliability guarantee where the recovery constraints in (4) and (5) are satisfied with a certain (high) probability.

3.1 Complexity Analysis

In this section, we investigate the complexity of WhiteRec. In fact, solving this problem involves finding a white channel assignment that is both *feasible* (in order to support the traffic demands as described in (1), (2), and (3)) and *optimal* (in order to minimize the recovery capacity as defined in (4) and (5)). Unfortunately, both of these problems are NP-hard.

Theorem 1. *Finding a feasible white channel assignment that sustains the given traffic demands is NP-hard.*

Proof. Our proof is based on the reduction from the Bin Packing Problem which is known to be NP-hard.

- Problem: Bin Packing Problem
- Input: a set of n items with volume $v_1, v_2, \ldots, v_n \in (0,1]$ and a set of m bins with unit capacity
- Decision: whether we can pack the n items into the m bins

To show the reduction, consider a star network with n links incident on a common node. The traffic demands on these links are $v_1, v_2, \ldots, v_n$. Suppose we have m homogeneous white channels, each with unit capacity. We would like to find a mapping from the n links to the m white channels such that the traffic demands are sustainable. Obviously, this is equivalent to determining the feasibility of packing the n items into the m bins. As a result, it is NP-hard to find a feasible assignment to support the given demands, even in a star network and when channels are homogeneous. □

Theorem 2. *Finding a white channel assignment that requires the minimum recovery capacity is NP-hard. Moreover, even if any channel assignment is feasible (i.e., the capacity of each white channel is sufficiently large such that the traffic demands are always sustainable under any channel assignment), the problem remains NP-hard.*

Proof. Our proof is based on the reduction from the Partition Problem which is known to be NP-hard [29, p. 223].

- Problem: Partition Problem
- Input: A set A of positive integers given by $A = \{r_1, \ldots, r_n\}$
- Output: A subset $S \subset A$ s.t. $\max\{sum(S), sum(A \setminus S)\}$ is minimized, where $sum(T)$ is the sum of all the elements in T

To show a mapping from the Partition Problem to our problem, consider a star network with n links, where link i has a traffic demand r_i. Suppose that we have two white channels w_1, w_2, and both channels have sufficiently large capacity such that any channel assignment is feasible. The goal is to find a white channel assignment requiring the minimum recovery channel capacity to recover from any single channel preemption. Hence, it is desirable to balance the loads on each white channel.

More formally, let $E(w)$ be the set of links using white channel w. It is easy to see that the recovery channel capacity in the formulation of WhiteRec can be expressed as

$$C = \max\left\{\sum\nolimits_{e\in E(w_1)} r_e, \sum\nolimits_{e\in E(w_2)} r_e\right\}.$$

In this setting, finding a white channel assignment minimizing C is equivalent to finding a subset S in the Partition Problem. This completes the proof. □

The above two theorems imply that WhiteRec is a complicated problem coupled by two NP-hard objectives: finding a feasible assignment to support the traffic demand and finding an optimal assignment that requires the minimum recovery capacity. To address this difficulty, we introduce a technique called *bi-criteria approximation* [30] which allows the feasibility constraints to be violated by a bounded amount while ensuring some approximation ratio with respect to the recovery capacity. The formal definition is as follows.

Definition 1 (Bi-Criteria Approximation). *An algorithm achieves (ρ, ϕ)-approximation to* WhiteRec *if the following two conditions are satisfied simultaneously.*

- *It requires at most ρ times of the minimum recovery capacity.*
- *It guarantees that at least ϕ-fraction of the traffic demand is sustained over each link.*

In the following sections, we first analyze the bi-criteria structure of WhiteRec. Based on the analysis, several approximation algorithms are developed and their bi-criteria approximation ratios are studied.

3.2 Bi-Criteria Analysis

In this section, we investigate the bi-criteria structure of the optimal feasible solution to WhiteRec. Specifically, we are interested in the structure that requires the minimum recovery capacity (i.e., optimality analysis, Section 3.2.1) and that sustains the given traffic demands (i.e., feasibility analysis, Section 3.2.2). Finally, the relationship between optimality and feasibility is discussed.

3.2.1 Optimality Analysis

We first study the structure of the optimal assignment that requires the minimum recovery capacity. The particular form of WhiteRec allows us to express the required recovery capacity C in a closed form. It is easy to see that constraints in (4) are equivalent to

$$C \geq M_1(y,k), \tag{6}$$

where

$$M_1(y,k) = \max_{v\in V, S\in \mathcal{W}(k)} \sum_{w\in S} \sum_{e\in\delta(v)} r_e y_e^w.$$

Similarly, constraints in (5) are equivalent to

$$C \geq M_2(y,k), \tag{7}$$

where

$$M_2(y,k) = \max_{U\in\mathcal{V}, S\in\mathcal{W}(k)} \frac{2}{|U|-1} \sum_{w\in S} \sum_{e\in E(U)} r_e y_e^w.$$

Combining (6) and (7), we can rewrite constraints (4) and (5) in WhiteRec as

$$C \geq \max\{M_1(y,k), M_2(y,k)\} \triangleq C(y,k). \tag{8}$$

In other words, given a white channel assignment y, the value of $C(y,k)$ is the minimum recovery capacity required to recover from any k channel preemptions. As a result, WhiteRec can be rewritten as

$$\begin{aligned} \min_y \quad & C(y,k) \\ \text{s.t.} \quad & y \text{ is feasible.} \end{aligned}$$

Note that $M_2(y,k)$ corresponds to the "odd-set constraints" in (5) which are difficult to handle in general. Hence, it is natural to consider the relaxation of WhiteRec by neglecting $M_2(y,k)$. The relaxed problem is referred to as WhiteRecApprox, i.e.,

$$\begin{aligned} \min_y \quad & M_1(y,k) \\ \text{s.t.} \quad & y \text{ is feasible.} \end{aligned}$$

The following lemma shows that the relaxation of $M_2(y,k)$ only leads to a small loss in optimality.

Lemma 1. *For any channel assignment y, we have*

$$M_1(y,k) \leq C(y,k) \leq 1.5 M_1(y,k). \tag{9}$$

Proof. The lower bound follows from the definition of $C(y,k)$. To show the upper bound, we notice that

$$
\begin{aligned}
M_2(y,k) &= \max_{U\in\mathcal{V},S\in\mathcal{W}(k)} \frac{2}{|U|-1}\sum_{w\in S}\sum_{e\in E(U)} r_e y_e^w \\
&= \frac{1}{2}\max_{U\in\mathcal{V},S\in\mathcal{W}(k)} \frac{2}{|U|-1}\sum_{w\in S}\sum_{v\in U}\sum_{e\in\delta(v)\cap E(U)} r_e y_e^w \\
&\le \frac{1}{2}\max_{U\in\mathcal{V},S\in\mathcal{W}(k)} \frac{2}{|U|-1}\sum_{v\in U}\sum_{w\in S}\sum_{e\in\delta(v)} r_e y_e^w \\
&\le \frac{1}{2}\max_{U\in\mathcal{V},S\in\mathcal{W}(k)} \frac{2}{|U|-1}\sum_{v\in U} M_1(y,k) \\
&= \frac{1}{2} M_1(y,k)\max_{U\in\mathcal{V}} \frac{2|U|}{|U|-1} = \frac{3}{2}M_1(y,k).
\end{aligned}
$$

The second inequality is due to the definition of $M_1(y,k)$, and the last equality holds because $|U|\ge 3$. □

Lemma 1 shows that the optimal solution to WhiteRecApprox attains 1.5-approximation to the original problem WhiteRec with respect to the required recovery capacity $C(y,k)$. In fact, in bipartite graphs, there is even no approximation gap between WhiteRecApprox and WhiteRec.

Theorem 3. *In a bipartite network, $M_1(y,k)\ge M_2(y,k)$ for any channel assignment y and any positive integer k.*

Proof. Let

$$
(U^*,S^*) = \max_{U\in\mathcal{V},S\in\mathcal{W}(k)} \frac{1}{\alpha(U)}\sum_{w\in S}\sum_{e\in E(U)} r_e y_e^w.
$$

The induced graph on U^* is denoted by $G^*=(U^*,E(U^*))$, which is still a bipartite graph. Let C be the minimum vertex cover of G^*. Then we have

$$
\begin{aligned}
M_2(y,k) &= \frac{1}{\alpha(U^*)}\sum_{e\in E(U^*)}\sum_{w\in S^*} r_e y_e^w \\
&\le \frac{1}{\alpha(U^*)}\sum_{v\in C}\sum_{e\in\delta(v)\cap E(U^*)}\sum_{w\in S^*} r_e y_e^w \\
&\le \frac{1}{\alpha(U^*)}\sum_{v\in C}\sum_{e\in\delta(v)}\sum_{w\in S^*} r_e y_e^w \\
&\le \frac{1}{\alpha(U^*)}\sum_{v\in C} M_1(y,k) = \frac{1}{\alpha(U^*)}|C|M_1(y,k),
\end{aligned}
$$

where the first inequality holds because every edge in $E(U^*)$ is incident on at least one node in C. By König's Theorem [31, pp. 203-207], the size of the minimum vertex cover equals to the size of the maximum matching in a bipartite graph. Thus, $|C|$ is upper-bounded by $\lfloor\frac{|U^*|}{2}\rfloor = \frac{|U^*|-1}{2}$ (note that $|U^*|$ is odd). Therefore, we can finally conclude that

$$
M_2(y,k) \le \frac{1}{\alpha(U^*)}|C|M_1(y,k) = M_1(y,k).
$$

□

This theorem implies that $C(y,k)=M_1(y,k)$ in a bipartite graph, and thus we can safely ignore $M_2(y,k)$ without sacrificing any optimality. In other words, the optimal solution to the relaxed problem WhiteRecApprox is also the optimal solution to the original problem WhiteRec if the secondary network is bipartite.

3.2.2 Feasibility Analysis

Next, we study the feasibility conditions (1), (2), and (3) and investigate the structure of channel assignments that are able to sustain the largest amount of traffic. In particular, we investigate the relationship between feasibility and optimality, which is important for our subsequent bi-criteria approximation analysis.

It is clear that deciding feasibility is equivalent to the following optimization problem FEASI which finds the maximum fraction of traffic that can be sustained over each link

$$
\begin{aligned}
\text{FEASI}:\quad & \max_{y,\beta}\ \beta \\
\text{s.t.}\ & \sum_{e\in\delta(v)} \frac{\beta r_e}{R_{w,e}} y_e^w \le 1, \forall v\in V, w\in W && (10)\\
& \sum_{e\in E(U)} \frac{\beta r_e}{R_{w,e}} y_e^w \le \frac{|U|-1}{2}, \forall U\in\mathcal{V}, w\in W && (11)\\
& \sum_{w\in W} y_e^w = 1, \forall e\in E \\
& y_e^w \in\{0,1\}, \forall e\in E, w\in W. && (12)
\end{aligned}
$$

Clearly, the original problem WhiteRec is feasible if and only if the optimal value β^* in FEASI is greater or equal to 1. Now let $\beta(y)$ be the maximum value of β in FEASI under an assignment y. The following lemma relates $\beta(y)$ to $C(y,1)$.

Lemma 2.

$$
\frac{R_{\min}}{C(y,1)} \le \beta(y) \le \frac{R_{\max}}{C(y,1)},
$$

where $R_{\min}$ and $R_{\max}$ are the minimum and the maximum white channel capacity, respectively.

Proof. From constraints (10) and (11), it follows that

$$
\beta(y) = \min\{Z_1(y), Z_2(y)\},
$$

where

$$
Z_1(y) = \max_{v\in V,w\in W} \frac{1}{\sum_{e\in\delta(v)} r_e y_e^w/R_{w,e}},
$$

$$
Z_2(y) = \min_{U\in\mathcal{V},w\in W} \frac{1}{\frac{2\sum_{e\in E(U)} r_e y_e^w/R_{w,e}}{|U|-1}}.
$$

Under the above notations, FEASI becomes

$$
\begin{aligned}
\max_{y}\ & \min\{Z_1(y), Z_2(y)\} \\
\text{s.t.}\ & \sum_{w\in W} y_e^w = 1, \forall e\in E \\
& y_e^w\in\{0,1\}, \forall e\in E, w\in W.
\end{aligned}
$$

It is clear that

$$
Z_1(y) \ge \frac{R_{\min}}{\max_{v\in V,w\in W}\sum_{e\in\delta(v)} r_e y_e^w} = \frac{R_{\min}}{M_1(y,1)},
$$

and similarly we have $Z_2(y)\ge \frac{R_{\min}}{M_2(y,1)}$. Then we obtain that

$$
\begin{aligned}
\beta(y) &\ge \min\left\{\frac{R_{\min}}{M_1(y,1)}, \frac{R_{\min}}{M_2(y,1)}\right\} \\
&= \frac{R_{\min}}{\max\{M_1(y,1), M_2(y,1)\}} = \frac{R_{\min}}{C(y,1)}.
\end{aligned}
$$

Similarly, we can show $\beta(y)\le\frac{R_{\max}}{C(y,1)}$. □

This lemma shows that if an assignment y yields a smaller $C(y,1)$, it tends to sustain more traffic. In particular, if white channels are homogeneous with capacity R, the lemma implies $\beta(y) = \frac{R}{C(y,1)}$. In this case, minimizing the recovery capacity required to survive a single preemption is equivalent to maximizing the amount of sustainable traffic. Therefore, this lemma bridges feasibility and optimality, which is important for our subsequent bi-criteria approximation analysis.

3.3 Algorithm 1: Greedy Algorithm

In this section, we propose a simple greedy algorithm to solve WhiteRec and analyze its bi-criteria approximation ratio. Without loss of generality, let the links in the secondary network be indexed by $e_1, \ldots, e_n$, where $n = |E|$. The greedy algorithm assigns a white channel to each of these links sequentially. Suppose we are deciding the channel assignment for link $e_i = (u, v)$, and define $\delta(u, v) = \delta(u) \cup \delta(v)$, i.e., the set of links incident on node u or node v. The greedy rule is to pick the white channel that currently sustains the least traffic over the links in $\delta(u, v)$. The detailed procedures are presented in Algorithm 1, where E_w corresponds to the set of links that are assigned channel w.

Algorithm 1. Greedy White Channel Assignment

1: Initialize $E_w = \emptyset,\ \forall w \in W$;
2: **for** $e_i = e_1, \ldots, e_n$ **do**
3: Assign white channel w^* to link $e_i = (u, v)$, where $w^* = \arg\min_{w \in W} \sum_{e \in \delta(u,v) \cap E_w} r_e$;
4: $E_{w^*} \leftarrow E_{w^*} \cup \{e_i\}$;
5: **end for**

The bi-criteria approximation ratio of this greedy algorithm is given in Theorem 4, where we define $R_{\min}$ and $R_{\max}$ to be the minimum and the maximum white channel capacity, respectively.

Theorem 4. *Suppose there exists a feasible solution to* WhiteRec. *Then the greedy algorithm achieves* $(\rho, \frac{1}{\rho}\frac{R_{\min}}{R_{\max}})$*-approximation to* WhiteRec, *where* $\rho = \frac{3}{2}(3 - \frac{2}{|W|})$.

Proof. See Section 3.6. □

For instance, if there are 2 homogeneous white channels, the greedy algorithm is guaranteed to sustain at least $\frac{1}{3}$ traffic demands while requiring less than 3 times of the minimum recovery capacity in WhiteRec.

The advantage of the greedy algorithm is in its simplicity. In fact, it does not require any global information when assigning channels for each individual link; thus, this greedy algorithm can even be implemented in a distributed manner, where more fresh local information can be used to improve the overall performance. Moreover, it is applicable to arbitrary networks. Although the theoretical approximation ratio of this algorithm is relatively loose, its practical performance turns out to be much better than the theoretical guarantee[1] (see Section 5). In fact, it is possible to improve the approximation ratio in a wide range of graphs. For example, it can be shown that this greedy algorithm achieves $(\frac{2}{3}\rho, \frac{3}{2\rho}\frac{R_{\min}}{R_{\max}})$-approximation to DetRec in a bipartite graph (by using Theorem 3).

1. Similar greedy algorithms have been shown to perform extremely well for frequency assignment in WDM-based optical networks [32].

3.4 Algorithm 2: Interference-Free Assignment

The above greedy algorithm is simple and has provable performance in any scenario but suffers from the relatively loose approximation ratio. In this section, we discuss an alternative channel assignment scheme, called Interference-Free Assignment (IFA), which is less general than the greedy algorithm but achieves much better performance.

Definition 2 (Interference-Free Assignment). *An assignment y is said to be interference-free if any two interfering links are assigned distinct white channels.*

For example, the channel assignment in Fig. 1b is interference-free while the one in Fig. 1a is not. Conceivably, IFA requires less recovery capacity since links that fail together due to any single channel preemption do not interfere with each other and can be activated simultaneously. Through the rest of this section, we study the properties of IFA. In particular, we will show IFA has near-optimal performance.

We first investigate the conditions for the existence of IFA. Note that IFA requires that adjacent links be assigned different channels; this is similar to *edge coloring* where each white channel corresponds to a color. From Vizing's Theorem [33] for edge coloring, we have the following observation:

Observation 1. There exists an interference-free channel assignment if the number of white channels is greater than the maximum node degree, i.e., $|W| > d_{\max}$.

The above observation shows that IFA does not always exist and is thus less general than the greedy algorithm. However, the condition shown in the above observation is very mild in practice since the number of white channels is usually much larger than the number of neighbors a node has [16].

Now we develop an algorithm for constructing an interference-free assignment (Algorithm 2). This algorithm gives an interference-free assignment whenever $|W| > d_{\max}$. Note that this algorithm is still valid if $|W| \leq d_{\max}$ but it does not have a provable performance in this case. Note also that this algorithm colors edges with white channels and there are several polynomial-time algorithms that can perform edge-coloring with $d_{\max} + 1$ colors in a simple graph (e.g., [34]), therefore Algorithm 2 can be run in polynomial time.

Algorithm 2. Interference-Free Channel Assignment

1: Color the graph with $d_{\max} + 1$ colors, which partitions the edges into $d_{\max} + 1$ matchings;
// These matchings are denoted by $I_1, \ldots, I_{d_{\max}+1}$.
2: **for** $i = 1 : d_{\max} + 1$ **do**
3: Assign edges in matching I_i to white channel w_i, where $w_i = i \bmod |W|$;
4: **end for**

Next, we investigate the properties of IFA. The most important one is given in Lemma 3 which shows that any interference-free channel assignment minimizes $M_1(y, k)$.

Lemma 3. *Consider any two interference-free channel assignments $\bar{y}$, $\tilde{y}$ and any non-interference-free assignment $\hat{y}$. Then the following relationship holds: $M_1(\bar{y}, k) = M_1(\tilde{y}, k) \leq M_1(\hat{y}, k)$ for all $k \in \mathbb{Z}^+$.*

Proof. For any interference-free assignment $\bar{y}$, let $\bar{S} \in \mathcal{W}(k)$ and $\bar{v} \in V$ be such that

$$M_1(\bar{y},k) = \sum_{w\in\bar{S}} \sum_{e\in\delta(\bar{v})} r_e \bar{y}_e^w. \tag{13}$$

Since $\bar{y}$ is interference-free, all the links incident on a node are assigned different white channels. This is also true for another interference-free assignment $\tilde{y}$. Thus, there exists a set $\tilde{S} \in \mathcal{W}(k)$ such that

$$\left\{ e \in \delta(\bar{v}) : \sum_{w\in\tilde{S}} \tilde{y}_e^w = 1 \right\} = \left\{ e \in \delta(\bar{v}) : \sum_{w\in\bar{S}} \bar{y}_e^w = 1 \right\}.$$

Therefore, we have

$$\sum_{w\in\bar{S}} \sum_{e\in\delta(\bar{v})} r_e \bar{y}_e^w = \sum_{w\in\tilde{S}} \sum_{e\in\delta(\bar{v})} r_e \tilde{y}_e^w,$$

which implies $M_1(\bar{y},k) \le M_1(\tilde{y},k)$ by the definition of $M_1(y,k)$. Similarly, we can prove $M_1(\bar{y},k) \ge M_1(\tilde{y},k)$. As a result, it follows that $M_1(\bar{y},k) = M_1(\tilde{y},k)$ for any interference-free channel assignments $\bar{y}$ and $\tilde{y}$.

To prove the second part, consider a non-interference-free channel assignment $\hat{y}$. Obviously, under the assignment $\hat{y}$, the preemption of k white channels can possibly lead to the preemption of more than k links incident on a node. Hence, there exists a set $\hat{S} \in \mathcal{W}(k)$ such that

$$\left\{ e \in \delta(\bar{v}) : \sum_{w\in\hat{S}} \hat{y}_e^w = 1 \right\} \supseteq \left\{ e \in \delta(\bar{v}) : \sum_{w\in\bar{S}} \bar{y}_e^w = 1 \right\}.$$

Therefore, we can conclude that $M_1(\bar{y},k) \le M_1(\hat{y},k)$. □

Lemma 3 together with Lemma 1 immediately implies that IFA achieves no more than 1.5 times the minimum recovery capacity. In fact, we can further tighten this bound, as shown in the following theorem.

Theorem 5. *Suppose there is a feasible solution to* WhiteRec *and an interference-free assignment exists. Then any interference-free assignment achieves* $(\frac{5}{4}, \frac{R_{\min}}{R_{\max}})$ *approximation to* WhiteRec.

Proof. We first prove that any IFA achieves no more than $\frac{5}{4}$ times the minimum recovery capacity. We start by introducing a lemma whose proof is similar to Lemma 1 and thus omitted. □

Lemma 4. *Let* $\mathcal{V}' = \{U \subseteq V : |U| \ge 5, |U| \text{ odd}\}$. *Then for any assignment* y *and integer* $k \ge 1$

$$\max_{U\in\mathcal{V}', S\in\mathcal{W}(k)} \frac{2}{|U|-1} \sum_{w\in S} \sum_{e\in E(U)} r_e y_e^w \le \frac{5}{4} M_1(y,k).$$

Now we get down to proving that any IFA achieves no more than $\frac{5}{4}$ times of the minimum recovery capacity. Denote $\mathcal{V}_3$ the collection of node sets with cardinality 3. For any channel assignment y and any integer $k \ge 1$, we rewrite $C(y,k)$ as

$$\begin{aligned} C(y,k) &= \max\{M_1(y,k), M_2(y,k)\} \\ &= \max\Bigg\{ M_1(y,k), \max_{U\in\mathcal{V}_3, S\in\mathcal{W}(k)} \frac{2}{|U|-1} \sum_{w\in S}\sum_{e\in E(U)} r_e y_e^w, \\ &\qquad \max_{U\in\mathcal{V}', S\in\mathcal{W}(k)} \frac{2}{|U|-1} \sum_{w\in S}\sum_{e\in E(U)} r_e y_e^w \Bigg\} \\ &\triangleq \max\{M_1(y,k), A(y,k), B(y,k)\} \end{aligned}$$

Let y' be an arbitrary IFA and y^* be the optimal solution to WhiteRec. We observe three key facts:

1) $A(y',k) \le A(y^*,k)$. This is due to the fact that in any induced graph of 3 nodes the interference-free assignment y' allocates different channels to different edges, which is optimal in that induced graph.
2) $B(y',k) \le \frac{5}{4} M_1(y',k) \le \frac{5}{4} M_1(y^*,k)$. This is due to Lemma 4 and the fact that any IFA minimizes $M_1(y,k)$.
3) $M_2(y',k) \ge M_1(y',k)$ otherwise $C(y',k) = M_1(y',k) \le M_1(y^*,k) \le C(y^*,k)$, which implies that y' is optimal. This fact further shows that $C(y',k) = \max\{A(y',k), B(y',k)\}$.

Then we have

$$\begin{aligned} \frac{C(y',k)}{C(y^*,k)} &= \frac{\max\{A(y',k), B(y',k)\}}{\max\{M_1(y^*,k), A(y^*,k), B(y^*,k)\}} \\ &\le \max\left\{ \frac{A(y',k)}{A(y^*,k)}, \frac{B(y',k)}{M_1(y^*,k)} \right\} \le \frac{5}{4}. \end{aligned}$$

Next, we prove that any IFA y' can sustain $\frac{R_{\min}}{R_{\max}}$-fraction of traffic over each link. Note that $C(y',1) = M_1(y',1)$ under the IFA y'. By Lemma 3, the IFA y' minimizes $M_1(y,1)$, so we have $C(y',1) = M_1(y',1) \le M_1(y,1) \le C(y,1)$ for any assignment y. Let $\hat{y}$ be a feasible solution to WhiteRec, i.e., $\beta(\hat{y}) \ge 1$. Then it follows that for any IFA y'

$$\beta(y') \ge \frac{R_{\min}}{C(y',1)} \ge \frac{R_{\min}}{R_{\max}} \frac{R_{\max}}{C(\hat{y},1)} \ge \frac{R_{\min}}{R_{\max}} \beta(\hat{y}) \ge \frac{R_{\min}}{R_{\max}},$$

where the first and third inequalities are due to Lemma 2, the second inequality is due to our claim that $C(y',1) \le C(\hat{y},1)$ and the last inequality holds because of our assumption that $\beta(\hat{y}) \ge 1$. This completes our proof.

Note that IFA has a much better approximation ratio than the greedy algorithm with respect to both the recovery capacity and the sustainable traffic. In particular, if channels are homogeneous, then any interference-free assignment is guaranteed to sustain 100 percent traffic demands while requiring less than 1.25 times the minimum recovery capacity. The caveat is that such a good approximation ratio only holds true when IFA exists (i.e., when $|W| > d_{\max}$).

In fact, IFA is even optimal with respect to the recovery capacity in many scenarios, as is shown in Theorem 6.

Theorem 6. *Suppose there is a feasible solution to* WhiteRec *and an interference-free assignment exists. Then any interference-free assignment achieves* $(1, \frac{R_{\min}}{R_{\max}})$ *approximation to* WhiteRec *in any of the following scenarios:*

i) $k = 1$, *i.e., we need to survive any single preemption;*
ii) $r_e = r\ \forall e \in E$ *and* $k \le d_{\max}$, *i.e., traffic demands are uniform and we want to survive no more than* $d_{\max}$ *preemptions;*
iii) *the secondary network is bipartite;*

Proof. We only prove that any IFA requires the minimum recovery capacity in these scenarios while the ratio for the sustainable traffic follows the same argument as in Theorem 5.

Part (i). Since y is interference-free and $k=1$, we have

$$M_1(y,k) = \max_{v\in V, w\in W} \sum_{e\in\delta(v)} r_e y_e^w = r_{\max},$$

$$M_2(y,k) \le r_{\max} \max_{U\in\mathcal{V}, w\in W} \frac{2}{|U|-1} \sum_{e\in E(U)} y_e^w$$
$$\le r_{\max} \max_{U\in\mathcal{V}, w\in W} \frac{2}{|U|-1} \left\lfloor \frac{|U|}{2} \right\rfloor = r_{\max}, \qquad (14)$$

where (14) is due to the fact that y is interference-free and consequently, the number of links in $U \subseteq V$ using the same channel is upper-bounded by the size of maximum matching in $E(U)$, which is $\lfloor \frac{|U|}{2} \rfloor$. The above two bounds show that $C(y,1) = r_{\max}$. Since $C(\tilde{y},1) \ge r_{\max}$ for any assignment $\tilde{y}$, we can conclude that y yields the minimum recovery capacity.

Part (ii). Consider an arbitrary interference-free channel assignment y. Since y is interference-free, all the links incident to a node are assigned different channels. Furthermore, we have $k \le d_{\max}$. Consequently, there exists a node v and $S \in \mathcal{W}(k)$ such that $\sum_{w\in S}\sum_{e\in\delta(v)} y_e^w = k$. Using this observation, the value $M_1(y,k)$ can be rewritten as

$$M_1(y,k) = r \max_{v\in V, S\in\mathcal{W}(k)} \sum_{w\in S}\sum_{e\in\delta(v)} y_e^w = rk.$$

On the other hand, using the similar trick to Part (i), we can bound $M_2(y,k)$ by $M_2(y,k) = rk$. Hence, we have $C(y,k) = M_1(y,k) = rk$. Clearly, $C(\tilde{y},k) \ge rk$ any assignment $\tilde{y}$; thus, y yields the minimum recovery capacity.

Part (iii). This directly follows from Theorem 3 and Lemma 3. □

The common feature of the above scenarios is that $C(y,k) = M_1(y,k)$ holds for any interference-free assignment y; as a result, Lemma 3 implies that any interference-free assignment minimizes $C(y,k)$ in these cases. Note that if white channels are homogeneous, then any interference-free assignment requires the minimum recovery capacity and 100 percent traffic demands can be sustained in any of the above scenarios. In other words, IFA is both feasible and optimal in these cases.

3.5 Scaling of Recovery Capacity

In this section, we investigate the scaling of the required recovery capacity under the proposed algorithms. Specifically, we show that the required recovery capacity becomes negligible as compared to the total traffic if the network size is relatively large.

To facilitate our analysis, we make a simplified assumption that traffic is uniform across the entire secondary network, i.e., $r_e = r$ for any $e \in E$. Also assume that white channels are homogeneous, i.e., $R_{w,e} = R$ for any $w \in W$ and $e \in E$. Denote by $C^*(k)$ the recovery capacity required to protect against any k channel preemptions under Algorithm 2. Also let L_{tot} be the total traffic demands in the secondary network, i.e., $L_{tot} = \sum_{e\in E} r_e = r|E|$. The following theorem shows the scaling of the relative recovery capacity ratio $C^*(k)L_{tot}$ with the network size $|V|$.

Theorem 7. *Suppose there is a feasible solution to* WhiteRec. *Then* $\frac{C^*(k)}{L_{tot}} = O(\frac{1}{|V|})$ *as* $|V| \to \infty$ *for any* $k \in \mathbb{Z}^+$.

Proof. Consider the channel assignment scheme shown in Algorithm 2. Clearly, each white channel is assigned to at most $\lceil \frac{d_{\max}+1}{|W|} \rceil$ matchings; thus, at most $\lceil \frac{d_{\max}+1}{|W|} \rceil$ links incident on the same node are assigned the same channel. Denote y the above white channel assignment. It follows that

$$\sum_{e\in\delta(v)} r_e y_e^w \le r \left\lceil \frac{d_{\max}+1}{|W|} \right\rceil, \forall w, v. \qquad (15)$$

Note that the matching sets derived in Algorithm 2 is also a matching set partition of $E(U)$ for each $U \in \mathcal{V}$. Hence, each white channel is assigned to at most $\lceil \frac{d_{\max}+1}{|W|} \rceil$ matchings in the matching set partition of $E(U)$. Since each matching of $E(U)$ has at most $\frac{|U|-1}{2}$ edges, it follows that for each $w \in W$ and $U \in \mathcal{V}$

$$\sum_{e\in E(U)} r_e y_e^w \le r \frac{|U|-1}{2} \left\lceil \frac{d_{\max}+1}{|W|} \right\rceil. \qquad (16)$$

By (15) and (16), we can see that $M_1(y,k)$ and $M_2(y,k)$ are upper bounded by

$$rk \left\lceil \frac{d_{\max}+1}{|W|} \right\rceil. \qquad (17)$$

It follows that

$$C^*(k) \le \max\{M_1(y,k), M_2(y,k)\}$$
$$\le rk \left\lceil \frac{d_{\max}+1}{|W|} \right\rceil.$$

If WhiteRec is feasible, then we have for any $v \in V$

$$\sum_{e\in\delta(v)}\sum_{w\in W} r_e y_e^w = r \sum_{e\in\delta(v)}\sum_{w\in W} y_e^w = r d_v \le |W|R,$$

where d_v is the degree of node v. Then it follows that $d_{\max} \le \frac{|W|R}{r}$, which implies that

$$C^*(k) \le rk \left\lceil \frac{R}{r} + \frac{1}{|W|} \right\rceil. \qquad (18)$$

At the same time, it is easy to see that

$$L_{tot} = r|E| \ge r\frac{d_{\min}|V|}{2} \ge \frac{r|V|}{2}. \qquad (19)$$

Dividing (18) by (19) yields the desired result. Note that R, r, k and $|W|$ are regarded as asymptotically constant factors when compared to $|V|$. □

Theorem 7 demonstrates that as the network size grows, the recovery capacity needed to protect against k white channel preemptions becomes negligible as compared to the total traffic, given that the recovery requirement k remains a constant. Our simulation results (see Section 5.2) show that the recovery capacity required to survive 2 preemptions is less than 1 percent of the total traffic in a 200-node network, even with very few white channels. This is mainly due to the effect of spatial reuse. That is, although the total traffic increases linearly with the network size, more links can be activated simultaneously; thus, the required recovery capacity does not scale up with the network size.

3.6 Proof of Theorem 4

Before the detailed proof, we first introduce a relaxed problem called WhiteRecInf, which is the same as WhiteRec but assumes infinite channel capacity such that any channel assignment can support the given traffic demands. In other words, feasibility conditions (1) and (2) are relaxed

$$\begin{aligned}\text{WhiteRecInf}: \quad & \min_{y} \quad C(y,k)\\ & \text{s.t.} \sum_{w\in W} y_e^w = 1, \forall e\in E\\ & y_e^w \in \{0,1\}, \forall e\in E, w\in W.\end{aligned}$$

We first show that the greedy algorithm yields no more than ρ times of the minimum recovery capacity in WhiteRec. We use the following notation:

- y_{OPT}: optimal solution to WhiteRec
- y^*: optimal solution to WhiteRecInf
- $\hat{y}$: solution given by the greedy algorithm
- $\hat{E}_w$: set of links that are assigned channel w under $\hat{y}$.

It is clear that $C(y^*,k) \le C(y_{\mathsf{OPT}},k)$ since WhiteRecInf is the relaxed problem of WhiteRec. Hence, it suffices to prove $C(\hat{y},k) \le \rho C(y^*,k)$. To facilitate our proof, we introduce a lemma.

Lemma 5. $C(y^*,k) \ge \frac{k}{|W|}\max_{v\in V}\sum_{e\in\delta(v)} r_e$.

Proof. Let v_1 be the node with the maximum traffic demands, i.e., $v_1 = \arg\max_{v\in V}\sum_{e\in\delta(v)} r_e$. Also denote L_w^* the total traffic supported by white channel w at node v_1 under assignment y^*. Without loss of generality, we assume $L_1^* \ge L_2^* \ge \cdots \ge L_{|W|}^*$. Then it follows that

$$M_1(y^*,k) = \sum_{1\le w\le k} L_w^*.$$

If $M_1(y^*,k) < \frac{k}{W}\sum_{e\in\delta(v_1)} r_e$, we would obtain

$$\frac{\sum_{1\le w\le k} L_w^*}{k} < \frac{1}{W}\sum_{e\in\delta(v_1)} r_e,$$

i.e., the average traffic (at node v_1) in the first k white channels are smaller than the average traffic (at node v_1) in all the white channels. This is an obvious contradiction since the first k channels support more loads at node v_1 than the remaining $|W|-k$ channels. Hence we can conclude that

$$C(y^*,k) \ge M_1(y^*,k) \ge \frac{k}{W}\sum_{e\in\delta(v_1)} r_e = \frac{k}{W}\max_{v\in V}\sum_{e\in\delta(v)} r_e.$$

This completes the proof to the lemma. □

Now we get back to proving $C(\hat{y},k) \le \rho C(y^*,k)$. Define

$$v^* = \arg\max_{v\in V, S\in\mathcal{W}(k)}\sum_{w\in S}\sum_{e\in\delta(v)\cap\hat{E}_w} r_e,$$

and without loss of generality, we suppose

$$\sum_{e\in\delta(v^*)\cap\hat{E}_1} r_e \ge \sum_{e\in\delta(v^*)\cap\hat{E}_2} r_e \ge \cdots \sum_{e\in\delta(v^*)\cap\hat{E}_{|W|}} r_e.$$

Note that under the above definitions, we have

$$M_1(\hat{y},k) = \sum_{1\le w\le k}\sum_{e\in\delta(v^*)\cap\hat{E}_w} r_e.$$

Suppose $e_w = (v^*, u_w^*)$ is the last edge added to $\hat{E}_w$ that is incident on v^*, and denote D_w the set of edges that have been assigned a white channel before edge e_w. Then it follows that for any white channel $w\in W$

$$\begin{aligned}\sum_{e\in\delta(v^*)\cap\hat{E}_w} r_e &= \sum_{e\in\delta(v^*)\cap\hat{E}_w\cap D_w} r_e + r_{e_w} && (20)\\ &\le \sum_{e\in\delta(v^*,u_w^*)\cap\hat{E}_w\cap D_w} r_e + r_{e_w} && (21)\\ &\le \frac{1}{|W|}\sum_{e\in\delta(v^*,u_w^*)\cap D_w} r_e + r_{e_w} && (22)\\ &\le \frac{1}{|W|}\Big(\sum_{e\in\delta(v^*)} r_e + \sum_{e\in\delta(u_w^*)} r_e\Big) + \frac{|W|-2}{|W|} r_{e_w} && (23)\\ &\le \frac{2}{|W|}\max_{v\in V}\sum_{e\in\delta(v)} r_e + \frac{|W|-2}{|W|} r_{e_w}. && (24)\end{aligned}$$

Here, (20) holds because edge e_w is the last one added to $\hat{E}_w$ that is incident on v^*; (22) is due to the fact edge e_w is assigned channel w only if channel w has the minimum aggregate loads at node v^* and u_w^* among all white channels (see step 1 in the greedy algorithm); (23) holds because e_w is incident on both v^* and u_w^* while D_w excludes e_w. Then we have

$$M_1(\hat{y},k) \le \frac{2k}{W}\max_{v\in V}\sum_{e\in\delta(v)} r_e + \frac{|W|-2}{|W|}\sum_{1\le w\le k} r_{e_w}.$$

By Lemma 5, we know

$$\frac{k}{|W|}\max_{v\in V}\sum_{e\in\delta(v)} r_e \le C(y^*,k).$$

At the same time, notice that $e_1, e_2, \ldots, e_{|W|}$ are distinct edges incident on v^*. Then it is easy to see that

$$\sum_{1\le w\le k} r_{e_w} \le C(y^*,k).$$

Therefore, we can conclude that

$$M_1(\hat{y},k) \le \left(3 - \frac{2}{|W|}\right)C(y^*,k) = \frac{2}{3}\rho C(y^*,k).$$

By Lemma 1, we finally have

$$C(\hat{y},k) \le \frac{3}{2}M_1(\hat{y},k) \le \rho C(y^*,k).$$

We now show that at least $\frac{1}{\rho}\frac{R_{\min}}{R_{\max}}$-fraction of traffic can be sustained by the greedy assignment. Let $\hat{y}$ be the solution obtained by the greedy algorithm, and denote by $\tilde{y}$ the optimal solution to FEASI. Then it follows from Lemma 2 that

$$\frac{\beta(\hat{y})}{\beta(\tilde{y})} \ge \frac{C(\tilde{y},1)}{C(\hat{y},1)}\frac{R_{\min}}{R_{\max}}. \qquad (25)$$

Denote OPT_1 the minimum recovery capacity required to survive one preemption in WhiteRec. Note that $\hat{y}$ is intended for surviving any k preemptions. However, the greedy algorithm is invariant to the number of preemptions we need to survive so $\hat{y}$ is also the greedy assignment for surviving one preemption. Thus, we have $C(\hat{y}, 1) \leq \rho\mathsf{OPT}_1$ according to the first approximation ratio. Also note that $C(\tilde{y}, 1) \geq \mathsf{OPT}_1$. Then

$$\frac{C(\tilde{y}, 1)}{C(\hat{y}, 1)} \geq \frac{\mathsf{OPT}_1}{\rho\mathsf{OPT}_1} = \frac{1}{\rho}. \tag{26}$$

Taking (26) into (25), we have

$$\frac{\beta(\hat{y})}{\beta(\tilde{y})} \geq \frac{1}{\rho}\frac{R_{\min}}{R_{\max}}.$$

Since WhiteRec has a feasible solution, then $\beta(\tilde{y}) \geq 1$ and $\beta(\hat{y}) \geq \frac{1}{\rho}\frac{R_{\min}}{R_{\max}}$. This completes the proof.

4 Probabilistic Recovery Requirement

Under the deterministic recovery requirement, the recovery capacity is provisioned to survive any k channel preemptions. However, some preemption events occur with a very low probability, and it is inefficient to protect against such rare events. In other words, some rare channel preemptions may be ignored without significantly affecting recovery performance. This motivates us to study probabilistic recovery requirements where the recovery capacity is provisioned in a way that the reliability is guaranteed with a high probability.

We assume that the availability of each white channel follows a two-state Markov chain, and a white channel is busy with probability p in the steady state. Since a channel should be preempted whenever it is busy, p is indeed the channel preemption probability if the secondary user is to use the channel. For each white channel w, define X_w to be a random variable which takes the value 1 if channel w is busy and 0 otherwise. Hence, X_w's are i.i.d. Bernoulli random variables with parameter p. Let $X = (X_1, \ldots, X_{|W|})$. Similar to the deterministic case, we define the following functions:

$$\tilde{M}_1(y, X) = \max_{v \in V} \sum_{w \in W} X_w \sum_{e \in \delta(v)} r_e y_e^w, \tag{27}$$

$$\tilde{M}_2(y, X) = \max_{U \in \mathcal{V}} \frac{2}{|U| - 1} \sum_{w \in W} X_w \sum_{e \in E(U)} r_e y_e^w, \tag{28}$$

$$\tilde{C}(y, X) = \max\{\tilde{M}_1(y, X), \tilde{M}_2(y, X)\}. \tag{29}$$

Clearly, $\tilde{C}(y, X)$ is a random variable, and its realization $\tilde{C}(y, x)$ is the recovery capacity needed to recover from the preemption scenario $X = x$. Note that $C(y, k)$ in (8) is the maximum of $\tilde{C}(y, x)$'s with $\sum_w x_w = k$. Also note that if the recovery capacity C is smaller than $\tilde{C}(y, x)$, it is impossible to fully recover from the preemption scenario $X = x$. Hence, our problem can be formulated as ProRec:

$$\begin{aligned} &\min_{C, y \text{ feasible}} \; C \\ &\text{s.t.} \;\; \Pr\left[\tilde{C}(y, X) > C\right] \leq \epsilon. \end{aligned} \tag{30}$$

Constraint (30) requires that any white channel preemption scenario should be survived with probability at least $1 - \epsilon$. This formulation allows us to ignore some rare preemption events while still meeting the probabilistic recovery requirement. Later in Section 5.4, we will numerically show that ignoring rare preemption events greatly reduces the recovery capacity. Finally, let $\tilde{C}(y, p, \epsilon)$ be the minimum recovery capacity required by channel assignment y for given p and ϵ.

Throughout this section, it is assumed that each white channel has sufficiently large capacity such that any assignment using one or more channels is feasible.[2] Under this assumption, we investigate how to find the optimal channel assignment that satisfies the probabilistic requirement with the minimum recovery capacity.

4.1 Probability Regimes

In the deterministic case, it is always better to use more white channels to survive any k preemptions, i.e., using more channels requires less recovery capacity. However, this is not always true in the case of probabilistic recovery requirements since using more white channels increases the probability that some white channels get preempted. In this section, we study various probability regimes and identify the optimal number of white channels to use. We use $|y|$ to denote the number of white channels used by assignment y.

The following theorem describes the scenario where the channel preemption probability p is relatively high.

Lemma 6. *Consider a white channel assignment y. We have $\tilde{C}(y, p, \epsilon) = C_{\max}$ if $p > \epsilon^{1/|y|}$, where $C_{\max}$ is the amount of recovery capacity needed to recover all the traffic.*

Proof. Note that $p^{|y|}$ is the probability that every white channel used by y gets preempted. Since $p^{|y|} > \epsilon$, we need to survive the scenario where every used channel fails and thus all the traffic in the network is lost. Otherwise, the recovery requirement (30) cannot be satisfied. Consequently, we conclude that $\tilde{C}(y, p, \epsilon) = C_{\max}$. □

Note that $C_{\max}$ is also an upper bound on the recovery capacity. This lemma shows that for each y, there exists a high probability regime where the required recovery capacity hits the upper bound. In other words, when p is relatively large, it is likely that we need to protect against all preemption scenarios. Based on this lemma, we characterize the search range for the optimal number of white channels to use in the high probability regime, given in Theorem 8.

Theorem 8. *Consider an arbitrary integer $K \in \{1, 2, \ldots, |W| - 1\}$. The optimal number of channels to use is in the range $\{K + 1, \ldots, |W|\}$ if $p > \epsilon^{\frac{1}{K}}$.*

Proof. Consider a channel assignment y that uses K channels. If $p > \epsilon^{\frac{1}{K}}$, then $\tilde{C}(y, p, \epsilon) = C_{\max}$ by Lemma 6. Moreover, for any channel assignment y' such that $|y'| < K$, we have $p > \epsilon^{\frac{1}{K}} > \epsilon^{\frac{1}{|y'|}}$. Again, it follows from Lemma 6 that $\tilde{C}(y', p, \epsilon) = C_{\max}$. Consequently, any channel assignment using K or fewer channels requires the maximum recovery capacity if $p > \epsilon^{\frac{1}{K}}$. This implies that removing those channel assignments from the set of feasible assignments does not affect the minimum recovery capacity needed to satisfy the recovery requirement. Thus, the optimal number of channels to use is $K + 1$ or more. □

Theorem 8 shows that it is preferable to use more white channels if the preemption probability is high. In contrast, it

2. This simplifying assumption allows us to focus on the optimal recovery capacity as opposed to feasibility. Moreover, the algorithm that will be presented in Section 4.2 does not rely on this assumption.

is better to use fewer white channels in the low probability regime, as is shown in the following theorem.

Theorem 9. *If there exists a positive integer K such that $p \le 1-(1-\epsilon)^{\frac{1}{K}}$, then any feasible white channel assignment y with $|y| \le K$ is optimal, requiring zero recovery capacity.*

Proof. Consider a channel assignment using $|y| \le K$ white channels. If $p \le 1-(1-\epsilon)^{\frac{1}{K}}$, we have $(1-p)^{|y|} \ge (1-p)^K \ge 1-\epsilon$, i.e., the probability that no preemption happens is greater than $1-\epsilon$ under assignment y. Thus, even if we do not provision any recovery capacity, the protection requirement is still satisfied. □

Theorem 9 describes a probability regime where p is so small and every preemption scenario occurs with probability less than ϵ. Thus, the probabilistic requirement is met even without using backup channels. In contrast, using more white channels will increase the probability that some white channel gets preempted, and non-zero recovery capacity may be needed in order to meet the probabilistic requirement. Note that the low preemption probability regime is largest when $K=1$, in which case the condition in Theorem 9 gives $p \le \epsilon$. Combining with Theorem 8, we can see that $p=\epsilon$ is the borderline between the high and the low probability regimes.

4.2 Probabilistic Recovery Algorithm

In this section, we present an algorithm to meet the probabilistic recovery requirement. The idea is to transform the intractable probabilistic requirement into a deterministic requirement that has been extensively studied in this paper.

Consider a feasible white channel assignment y that uses $|y|$ white channels. For a nonnegative integer k, let $B(|y|,k)$ be the probability that more than k white channels (out of the $|y|$ used channels) are preempted, i.e.,

$$B(|y|,k) = \sum_{i=k+1}^{|y|} \binom{|y|}{i} p^i (1-p)^{|y|-i}.$$

Clearly, if $B(|y|,k) \le \epsilon$, it is unnecessary to protect against more than k preemptions. However, when $B(|y|,k) > \epsilon$, it is necessary to protect against k channel preemptions *in a deterministic way*. To reduce the required recovery capacity, it is natural to consider the smallest k that satisfies $B(|y|,k) \le \epsilon$, and such a k is denoted by $\Gamma(|y|,p,\epsilon)$. Now, in order to satisfy the probabilistic requirement, it is necessary to protect against all preemption events with fewer than (or equal to) $\Gamma(|y|,p,\epsilon)$ preemptions. Note that computing $\Gamma(|y|,p,\epsilon)$ requires the value of $|y|$, where we can leverage our previous analysis about probabilistic regimes to find the optimal number of white channels to use (this will be elaborated on shortly).

Now we have all the ingredients for the algorithm. First, we compare the values of p and ϵ to determine the probability regime.

- If $p \le \epsilon$, we look at the low probability regime and find the largest K (say K_0) such that $p \le 1-(1-\epsilon)^{\frac{1}{K}}$. By Theorem 9, we know that using K_0 white channels is optimal. Note that we choose the largest K in order to sustain more traffic demands. Then we select K_0 white channels with the largest capacity. The remaining task is to find a feasible assignment using these channels, which has been discussed in the deterministic setting.
- If $p > \epsilon$, we look at the high probability regime and find the largest K (say K_0) such that $p > \epsilon^{\frac{1}{K}}$. By Theorem 8, the optimal number of white channels to use can only be K_0+1 or more. Hence, we need to search in the set $\{K_0+1,\ldots,|W|\}$, and the search procedures are as follows. For each $J \in \{K_0+1,\ldots,|W|\}$, we compute $\Gamma(J,p,\epsilon)$ and select J channels with the largest capacity. Then we solve a deterministic recovery problem (i.e., WhiteRec) where there are J channels and we need to survive $k=\Gamma(J,p,\epsilon)$ preemptions. Note that each J will have a corresponding recovery capacity $C^*(J)$ and assignment $y^*(J)$. The last task is to select a J with the minimum $C^*(J)$ and return the corresponding assignment.

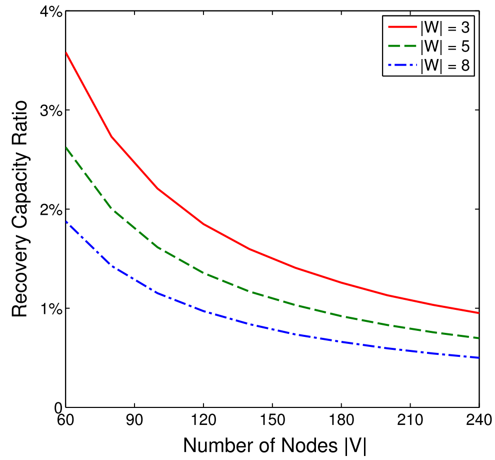

Fig. 2. Scaling of the relative recovery capacity ratio with the network size $|V|$ (where $k=2$ channel preemptions are to be survived).

5 Performance Evaluation

In this section, we numerically study our schemes. Specifically, we seek to answer the following questions:

- How does the recovery capacity scale with the network size?
- What is the bi-criteria approximation quality of the greedy algorithm and IFA?

5.1 Simulation Setup

We use Erdős-Renyi Random Graph to simulate the network topology, where links are established with probability 0.6 and the maximum node degree is bounded by 8. The traffic demand over each link is uniformly distributed in the range [1,100] Mbps. The capacity of each white channel is uniformly distributed in the range [75,200] Mbps. In our simulation, 5,000 random graph instances are tested.

5.2 Scaling of Recovery Capacity

We first investigate how the relative recovery capacity ratio (see Section 3.5 for the definition) scales with the network size. As is observed in Fig. 2, the recovery capacity ratio goes down with the growth of the network size. Specifically, the required recovery capacity is only around 1 percent of the total traffic demands in a 200-node network, even with very few white channels (e.g., $|W|=3$). Therefore, we expect the recovery capacity to become negligible as compared to the total traffic

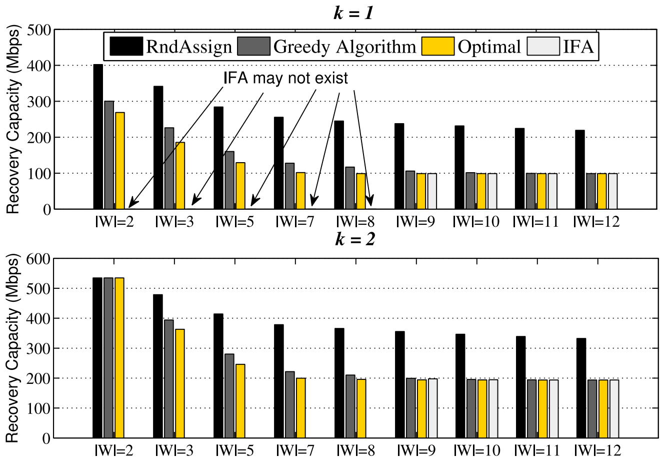

Fig. 3. The comparison among different algorithms with respect to the required recovery capacity ($|V| = 20$).

demands as the network size continues to grow. In addition, curve fitting shows that the recovery capacity ratio scales as $\Theta(\frac{1}{|V|^a})$ where a ranges in 1.02-1.09, which roughly matches the theoretical bound we obtain in Theorem 7. As mentioned in Section 3.5, this result is valid when k is a constant.

5.3 Approximation Quality

Since we consider the bi-criteria approximation framework, two metrics should be evaluated: the recovery capacity and the fraction of traffic sustained over each link. Through the rest of this section, we study the two aspects by comparing the following schemes.

- Greedy Algorithm (Algorithm 1).
- Interference-Free Assignment (IFA). Note that an interference-free assignment is guaranteed to exist only if $|W| > d_{\max}$ (in our simulation, $d_{\max} = 8$).
- Random Assignment (RndAssign) that assigns each link a random white channel.
- Optimal result to WhiteRec, computed using Gurobi, a large-scale mathematical programming solver.

5.3.1 Recovery Capacity

Fig. 3 illustrates the comparison of these schemes with respect to the recovery capacity; Table 1 lists the detailed approximation gap.[3] We first focus on the approximation quality of IFA. When $k = 1$, IFA yields the same amount of recovery capacity as the optimal solution and the approximation gap is zero. Recall that IFA is optimal when $k = 1$ (see Theorem 6). When $k = 2$, IFA is only slightly worse than the optimum (less than 2 percent, as is shown in Table 1), much better than the 1.25-approximation bound. The only caveat is that IFA is guaranteed to exist only if $|W| > d_{\max}$. Note that IFA exists when the number of white channels is greater than the maximum node degree, which is bounded by 8 in our simulations. As Theorem 6 suggests, in many scenarios, IFA is optimal if one exists. This is why IFA is close to optimal when $|W| > 8$.

Next, we investigate the approximation quality of the greedy algorithm. Despite its relatively loose approximation ratio, the greedy algorithm performs very well in practice.

3. The approximation gap is defined by $\frac{\text{ALG}-\text{OPT}}{\text{OPT}}$, where ALG is the amount of required recovery capacity by using the approximation algorithm and OPT is the minimum recovery capacity.

TABLE 1
Approximation Gap of Different Schemes

	Survive $k = 1$ failures			Survive $k = 2$ failures		
$\|W\|$	Rnd	Greedy	IFA	Rnd	Greedy	IFA
2	50%	12%	N/A	0%	0%	N/A
3	84%	22%	N/A	32%	9%	N/A
5	120%	24%	N/A	68%	14%	N/A
7	151%	26%	N/A	90%	11%	N/A
8	148%	18%	N/A	87%	7%	N/A
9	140%	7%	0%	83%	3%	2%
10	134%	3%	0%	79%	1%	0%
11	127%	1%	0%	75%	0%	0%
12	122%	0%	0%	72%	0%	0%

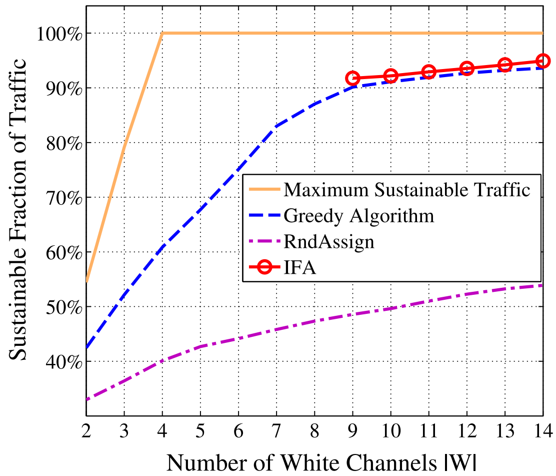

Fig. 4. Comparison among different algorithms with respect to the fraction of traffic sustained over each link ($|V| = 20$).

The worst approximation gap is 26 percent when $k = 1$ and 14 percent when $k = 2$. It also outperforms the random assignment by almost an order of magnitude in terms of the approximation gap. For a given link, the greedy algorithm seeks to find a channel that has not been assigned to the neighboring links or has been assigned to the least (in total load) number of neighboring links. Hence, our greedy algorithm prefers an interference-free channel assignment. This is why the greedy algorithm performs comparably to IFA when the number of channels is greater than 8 (when IFA always exists). Note, however, that the greedy algorithm sequentially selects a link and assign a channel without the knowledge of network-wide assignment, it is not guaranteed to find IFA even if one exists. Although the greedy algorithm performs slightly worse than IFA, it has the advantage of being applicable in any scenario.

5.3.2 Sustainable Traffic

In Fig. 4, we illustrate the comparison among different assignment schemes in terms of the fraction of traffic sustained over each link. Note that the maximum sustainable traffic level is obtained by solving FEASI (see Section 3.2.2) in Gurobi. We first notice that if there is only a small number of white channels, the maximum sustainable traffic level can be less than 100 percent. With more white channels, we have more spectrum resources and 100 percent traffic demands are sustainable. By comparison, the fraction of traffic sustained by the greedy algorithm is reasonably good

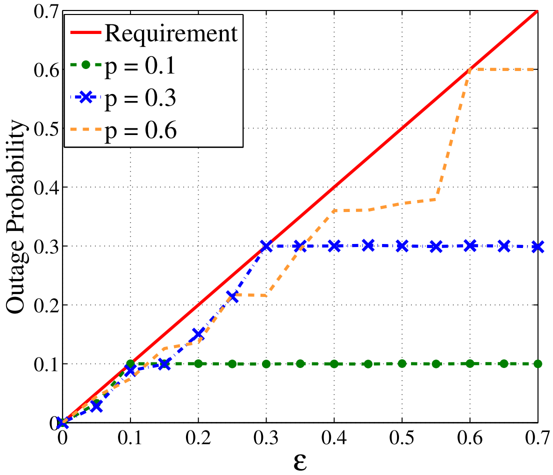

Fig. 5. The probability that the secondary network is not survivable against channel preemptions ($|V| = 20$).

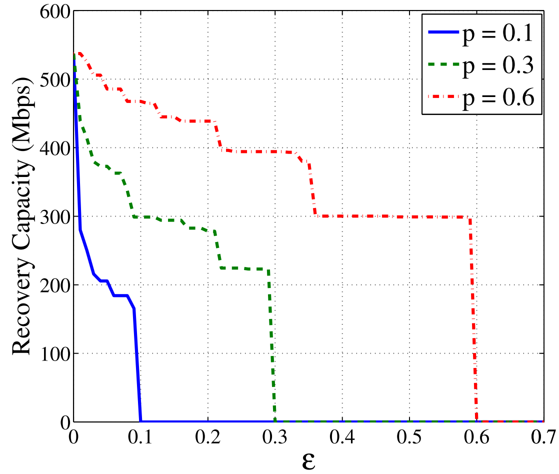

Fig. 6. The impact of ϵ on the required recovery capacity ($|V| = 20$).

as compared to the maximum sustainable level (at least 60 percent of the maximum), and the greedy algorithm significantly outperforms the random assignment. In particular, given a sufficient number of white channels (say $|W| \geq 9$), the greedy algorithm yields comparable performance to IFA and sustains over 90 percent traffic demands.

5.4 Probabilistic Recovery Requirement

In the probabilistic setting, the availability of each white channel follows a two-state Markov chain with the steady-state busy probability p. We need to survive any preemption scenario with probability $1 - \epsilon$.

5.4.1 Outage Probability

First, we need to evaluate the outage probability yielded by our probabilistic recovery algorithm, i.e., the probability that the secondary network does not recover from channel preemptions. To estimate such a probability, we run the Monte Carlo method with 10,000 samples, and Fig. 5 shows the result. We can observe that the outage probability yielded by our algorithm is always lower than the probabilistic recovery requirement ϵ, thus validating the legitimacy of our algorithm.

5.4.2 Impact of ϵ

Next, we discuss the influence of ϵ on the amount of recovery capacity we need provision. As can be seen from Fig. 6, if white channels are preempted with low probability, even a very small value of ϵ can significantly reduce the required recovery capacity. For example, if $p = 0.1$, we can reduce the required recovery capacity by almost 50 percent if we require 99 percent rather then 100 percent reliability. On the other hand, if the channel preemption probability is high (say $p = 0.6$), we do not observe much reduction in the recovery capacity, unless we sacrifice a significant degree of reliability (i.e., ϵ needs to be relatively large to significantly reduce the recovery capacity). In fact, with a high channel preemption probability, it is very likely that a large number of white channels are preempted and we need to protect against almost all preemption events. Therefore, probabilistic recovery requirements are particularly beneficial in the case where channel preemptions are rare events.

6 Conclusions

In this paper, we use backup channels to provide reliability guarantees for secondary users. In particular, we investigate the optimal white channel assignment that minimizes the recovery capacity required to meet either a deterministic or a probabilistic recovery requirement. This problem is shown to be coupled by two NP-hard objectives, and several bi-criteria approximation schemes are developed. For probabilistic recovery requirements, we also characterize different probability regimes and develop a heuristic algorithm. In this work, we assumed that the traffic demand on every link is fixed. However, in practice, the traffic demands are dynamic as there are arrivals and departures of service requests. When there is a new arrival of traffic, it may be necessary to find a route from source to destination, and even update the channel assignment; so that reliability requirements can be fulfilled. Hence, it would be interesting to incorporate dynamic arrivals of traffic demands into our framework and solve the routing and channel assignment problem. We leave this problem as future study.

Acknowledgments

This work was presented in part at WiOpt 2016 [1]. This work was supported by US National Science Foundation Grants CNS-1116209, CNS-1524317, CNS-1701964, and AST-1547331. The work of H.-W. Lee was supported by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (No.2018R1D1A1B07048388).

References

[1] Q. Liang, H.-W. Lee, and E. Modiano, "Robust design of spectrum-sharing networks," in *Proc. 14th Int. Symp. Model. Optimization Mobile Ad Hoc Wireless Netw.*, May 2016, pp. 1–8.
[2] J. Mitola III, "Cognitive radio for flexible mobile multimedia communications," *ACM/Kluwer Mobile Netw. Appl.*, vol. 6, no. 5, pp. 435–441, Sep. 2001.
[3] FCC Order DA 11-131, "Unlicensed Operation in the TV Broadcast Bands; Additional Spectrum for Unlicensed Devices Below 900 MHz and in the 3 GHz Band," Jan. 2011.
[4] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: A POMDP framework," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 3, pp. 589–600, Apr. 2007.
[5] A. C. C. Hsu, D. S. L. Wei, and C. C. J. Kuo, "A cognitive MAC protocol using statistical channel allocation for wireless ad-hoc networks," in *Proc. IEEE Wireless Commun. Netw. Conf.*, Mar. 2007, pp. 105–110.

[6] W. Y. Lee and I. F. Akyldiz, "A spectrum decision framework for cognitive radio networks," *IEEE Trans. Mobile Comput.*, vol. 10, no. 2, pp. 161–174, Feb. 2011.

[7] K. G. M. Thilina, E. Hossain, and D. I. Kim, "DCCC-MAC: A dynamic common-control-channel-based MAC protocol for cellular cognitive radio networks," *IEEE Trans. Veh. Technol.*, vol. 65, no. 5, pp. 3597–3613, May 2016.

[8] I. Kakalou, G. I. Papadimitriou, P. Nicopolitidis, P. G. Sarigiannidis, and M. S. Obaidat, "A reinforcement learning-based cognitive MAC protocol," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2015, pp. 5608–5613.

[9] H. Kim and K. G. Shin, "Fast discovery of spectrum opportunities in cognitive radio networks," in *Proc. 3rd IEEE Symp. New Frontiers Dyn. Spectr. Access Netw.*, Oct. 2008, pp. 1–12.

[10] Q. Liang, S. Han, F. Yang, G. Sun, and X. Wang, "A distributed-centralized scheme for short- and long-term spectrum sharing with a random leader in cognitive radio networks," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 11, pp. 2274–2284, Dec. 2012.

[11] S. Sengottuvelan, J. Ansari, P. Mhnen, T. G. Venkatesh, and M. Petrova, "Channel selection algorithm for cognitive radio networks with heavy-tailed idle times," *IEEE Trans. Mobile Comput.*, vol. 16, no. 5, pp. 1258–1271, May 2017.

[12] K. F. Li, W. C. Lau, and O. C. Yue, "Link restoration in cognitive radio networks," in *Proc. IEEE Int. Conf. Commun.*, May 2008, pp. 371–376.

[13] K. R. Chowdhury and I. F. Akyildiz, "Cognitive wireless mesh networks with dynamic spectrum access," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 1, pp. 168–181, Jan. 2008.

[14] I. F. Akyildiz, W.-Y. Lee, and K. R. Chowdhury, "CRAHNs: Cognitive radio ad hoc networks," *Ad Hoc Netw.*, vol. 7, no. 5, pp. 810–836, 2009.

[15] M. Youssef, M. Ibrahim, M. Abdelatif, L. Chen, and A. V. Vasilakos, "Routing metrics of cognitive radio networks: A survey," *IEEE Commun. Surveys Tuts.*, vol. 16, no. 1, pp. 92–109, Jan. 2014.

[16] Q. Liang, X. Wang, X. Tian, F. Wu, and Q. Zhang, "Two-dimensional route switching in cognitive radio networks: A game-theoretical framework," *IEEE/ACM Trans. Netw.*, vol. 23, no. 4, pp. 1053–1066, Aug. 2015.

[17] P. K. Tseng and W. H. Chung, "Local rerouting and channel recovery for robust multi-hop cognitive radio networks," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2013, pp. 2895–2899.

[18] J. Zhao and G. Cao, "Robust topology control in multi-hop cognitive radio networks," *IEEE Trans. Mobile Comput.*, vol. 13, no. 11, pp. 2634–2647, Nov. 2014.

[19] X. Wang, M. Sheng, D. Zhai, J. Li, G. Mao, and Y. Zhang, "Achieving Bi-channel-connectivity with topology control in cognitive radio networks," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 11, pp. 2163–2176, Nov. 2014.

[20] M. Cardei and A. Mihnea, "Channel assignment in cognitive wireless sensor networks," in *Proc. Int. Conf. Comput. Netw. Commun.*, Feb. 2014, pp. 588–593.

[21] I. A. M. Balapuwaduge, F. Y. Li, A. Rajanna, and M. Kaveh, "Channel occupancy-based dynamic spectrum leasing in multichannel CRNs: Strategies and performance evaluation," *IEEE Trans. Commun.*, vol. 64, no. 3, pp. 1313–1328, Mar. 2016.

[22] I. A. M. Balapuwaduge, F. Y. Li, and V. Pla, "Dynamic spectrum reservation for CR networks in the presence of channel failures: Channel allocation and reliability analysis," *IEEE Trans. Wireless Commun.*, vol. 17, no. 2, pp. 882–898, Feb. 2018.

[23] K. Jagannathan, I. Menache, E. Modiano, and G. Zussman, "Noncooperative spectrum access; the dedicated versus free spectrum choice," *IEEE J. Sel. Areas Commun.*, vol. 30, no. 11, pp. 2251–2261, Dec. 2012.

[24] E. Modiano, D. Shah, and G. Zussman, "Maximizing throughput in wireless networks via gossiping," in *Proc. Joint Int. Conf. Meas. Model. Comput. Syst.*, 2006, pp. 27–38.

[25] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle, "Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks," in *Proc. IEEE INFOCOM*, Apr. 2006, pp. 1–13.

[26] S. Singh, R. Mudumbai, and U. Madhow, "Distributed coordination with deaf neighbors: Efficient medium access for 60 GHz mesh networks," in *Proc. IEEE INFOCOM*, Mar. 2010, pp. 1–9.

[27] J. Edmonds, "Maximum matching and a polyhedron with (0, 1) vertices," *J. Res. Nat. Bureau Standards*, vol. 69, pp. 125–130, 1965.

[28] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," *IEEE Trans. Inf. Theory*, vol. 34, no. 5, pp. 910–917, Sep. 1988.

[29] M. R. Garey and D. S. Johnson, *Computers and Intractability; A Guide to the Theory of NP-Completeness*. New York, NY, USA: Freeman, 1990.

[30] M. Ehrgott, *Multicriteria Optimization*. Berlin, Germany: Springer, 2005.

[31] N. Biggs, E. Lloyd, and R. Wilson, *Graph Theory 17361936*. Oxford, U.K.: Clarendon, 1999.

[32] S. Subramaniam and R. A. Barry, "Wavelength assignment in fixed routing WDM networks," in *Proc. IEEE Int. Conf. Commun.*, Jun. 1997, pp. 406–410.

[33] V. Vizing, "On an estimate of the chromatic class of a p-graph," *Diskret. Analiz.*, vol. 3, pp. 25–30, 1964.

[34] J. Misra and D. Gries, "A constructive proof of Vizing's theorem," *Inf. Process. Lett.*, vol. 41, no. 3, pp. 131–133, 1992.

Qingkai Liang received the BE degree (with the highest honor) from Shanghai Jiao Tong University, China, in 2013 and the MS degree from the Massachusetts Institute of Technology (MIT), in 2015. He is currently working toward the PhD degree at the Massachusetts Institute of Technology. His research interests include various learning and control problems that arise in networked systems such as data center networks, wireless networks, and optical networks. He is a student member of the IEEE.

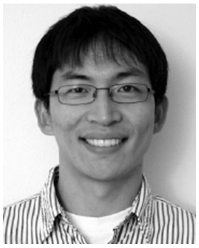

Hyang-Won Lee (S'02-M'08) received the BS, MS, and PhD degrees in all electrical engineering and computer science from the Korea Advanced Institute of Science and Technology, Daejeon, Korea, in 2001, 2003, and 2007, respectively. He was a postdoctoral research associate with the Massachusetts Institute of Technology during 2007-2011, and a research assistant professor with KAIST between 2011 and 2012. He is currently an associate professor with the Department of Software, Konkuk University, Seoul, Republic of Korea. His research interests include the areas of network algorithms, robust optimization, and reliable network design. He is a member of the IEEE.

Eytan Modiano received the BS degree in electrical engineering and computer science from the University of Connecticut, Storrs, in 1986, and the MS and PhD degrees, both in electrical engineering, from the University of Maryland, College Park, Maryland, in 1989 and 1992, respectively. He is a professor with the Department of Aeronautics and Astronautics and associate director of the Laboratory for Information and Decision Systems (LIDS), MIT. Prior to joining the faculty with MIT in 1999, he was a Naval Research Laboratory fellow between 1987 and 1992, a National Research Council post doctoral fellow during 1992-1993, and a member of the technical staff with MIT Lincoln Laboratory between 1993 and 1999. His research is on modeling, analysis, and design of communication networks and protocols. He is the co-recipient of the INFOCOM 2018 Best Paper Award, the MobiHoc 2018 Best Paper Award, MobiHoc 2016 Best Paper Award, the Wiopt 2013 Best Paper Award, and the SIGMETRICS 2006 Best Paper Award. He is the editor-in-chief of the *IEEE/ACM Transactions on Networking*, and served as an associate editor of the *IEEE Transactions on Information Theory* and the *IEEE/ACM Transactions on Networking*. He was the technical program co-chair for IEEE Wiopt 2006, IEEE Infocom 2007, ACM MobiHoc 2007, and DRCN 2015. He served on the IEEE Fellows Committee in 2014 and 2015, and is a fellow of the IEEE and an associate fellow of the AIAA.