Centralized, Measurement-based, Spectrum Management for Environments with Heterogeneous Wireless Networks

George Nychis, Charalampos E. Tsourakakis, Srinivasan Seshan, Peter Steenkiste
Carnegie Mellon University

***Abstract*—Heterogeneity of wireless networks has become an increasing problem in the wireless spectrum that breaks down spectrum sharing and exacerbates interference. Many coexistence techniques have been proposed to alleviate this interference, however, they are difficult to deploy due to changes needed in the protocols, overhead, and rapid changes in technology.**

In this paper, we focus on the potential of spectrum management to provide a long-term solution. We introduce novel components to a spectrum management system that overcomes limitations of current models that have remained relatively focused on homogeneous environments. Our approach is a centralized one, where we analyze information collected from heterogeneous monitors available today, structure the information in a hypergraph, and perform an analysis to detect heterogeneous conflicts. Introducing a mixed integer program (in addition to other novel components), we reconfigure devices in the spectrum to avoid conflicts and improve performance.

I. INTRODUCTION

As spectrum use becomes increasingly heterogeneous, wireless network performance can suffer greatly due to interference driven by diversity between the many protocols and standards in the spectrum [9], [16], [26], [27]. Given the fundamental inability for a single protocol or radio to meet the specific needs and constraints of all applications for wireless technology, diversity and increased interference are trends that are likely to continue over time.

In an attempt to address diversity and interference, significant effort has been focused on developing coexistence techniques between technologies (e.g., [1], [9], [14], [16], [21], [32]). These techniques introduce modifications to the radios and protocols that reduces interference between two specific technologies when operating in the same band. While coexistence techniques can alleviate interference, this general approach requires N^2 solutions between all technology pairs (where N is growing). Additionally, they are difficult to deploy. They incur overhead, changes are often needed at lower layers (e.g., PHY & MAC), and rapid changes to technologies and standards can make such solutions short-lived.

To the contrary, spectrum management has the ability to provide a long-term and "single" solution through a system that frequency-isolates incompatible technologies when possible, and otherwise intelligently places them together where they will receive and generate the least interference. Unfortunately, current approaches in spectrum management have remained fairly homogeneous. For example, many works assume all networks within range will coordinate and that all networks have the same potential sets of center frequencies and bandwidths (i.e., channels) [2], [22], [23], [28]. More recent work that considers heterogeneity between networks unfortunately remains Wi-Fi centric [26], i.e., it only looks to optimize a Wi-Fi network to avoid heterogeneous interference.

In this paper, we focus on overcoming such limitations of prior work in spectrum management systems, and explore spectrum management's potential at reducing interference between heterogeneous networks without the need for N^2 coexistence protocols. In particular, we focus on small to moderate size environments where centralized and measurement-based management is possible, and on limitations across 3 common components of spectrum management systems: 1) The RF environmental model that represent the networks, devices, and their signal characteristics in an environment, 2) The spectrum assignment algorithm that determines the frequencies for each network using the RF environmental model (or, a conflict graph derived from it), and 3) Predictive channel quality metrics that estimate the performance of a network on a specific channel to predict the performance of various configurations.

We present novel contributions in all 3 of these key areas, leading to effective and efficient spectrum management for heterogeneous networks. Our system design is *not* Wi-Fi centric, and it follows design principles that we introduce, meant to keep the system as general and generic as possible to support various technologies and their evolution over time.

Our contributions are as follows. First, we introduce a novel *hypergraph-based RF environmental model* that is able to represent rich information about the environment and differences between the heterogeneous networks and technologies in it. A major benefit of our graph-based model is the ability to flexibly search it for various (but specific) relationships between its components (i.e., networks and radios). Leveraging this flexibility, we introduce subgraph templates that represent various types of conflicts and apply subgraph isomorphism (or, subgraph matching) to detect each conflict in the environment.

Second, we present a new predictive channel quality metric that considers heterogeneity between networks and devices. The metric estimates the expected airtime of a radio on a particular channel by: 1) Accounting for its fair share of airtime from networks it coordinates with, and 2) Degrading this expected airtime due to interference from heterogeneous networks. The degradation is calculated using fundamental properties of the radios such as their airtimes, and whether both radios are unable to coordinate with each other, or whether at least one is able to coordinate (i.e., an asymmetric scenario).

Finally, we provide the hypergraph-based model of the environment to a *mixed-integer program (MIP) based optimization*. The algorithm uses the constraints given in the model (e.g., the possible frequencies of each radio), decomposes the hypergraph in to a series of conflicts (similar to a conflict graph), and uses our predictive metric to find efficient organizations that reduce interference from heterogeneous networks.

Through a real world evaluation with heterogeneous networks, we show that our spectrum management system (i.e., a "single" solution) can detect and avoid various types of conflicts between heterogeneous networks reported in prior works. Additionally, we compare our assignments to the current state of the art in today's environments: first-come, first-served. Even with heavily constrained spectrum (e.g., where it is not possible to completely frequency isolate heterogeneous networks), our system can reduce loss and improve performance.

The remainder of this paper is organized as follows. We present the appropriate background with limitations of current systems in Section II. The principles of our design in Section III, followed by our system and its components in Section IV. We present our evaluation in Section V, and conclude with discussion in Section VI.

II. BACKGROUND AND CURRENT PRACTICE

A. *Spectrum Management and its Key Components*

The goal of spectrum management is to (re)organizing networks in the frequency domain to achieve some objective function. For example, to minimize the overlap of networks in the spatial and frequency domains [11], or to prioritize frequency assignments based on the traffic loads of networks in the spectrum [22], [23], [28]. The goal of our work is to organize the spectrum in a way that improves performance by minimizing interference between heterogeneous networks.

From studying prior work (e.g., [7], [18], [25], [28]), we have found that many spectrum management systems have a similar structure with **3 main components**, shown in Figure 1. Above the underlying monitoring infrastructure:

1) An *RF environmental model* that takes the raw information from the monitor, and provides structure to it. This model typically provides constraints in the system, as well as meaningful information about the interactions of the devices through analysis. Examples include a conflict graph that represents radios that may interfere [13], [17], [31], a weighted graph to convey traffic load and spatial overlap [18], [20], network topologies using cliques, and graph-based topology including interference and coverage ranges [5], [30].

2) An *assignment algorithm* evaluates different frequency-domain configurations of the networks and devices, given the constraints and information provided by the environmental model. The goal of the algorithm is to find potential configurations that reduce conflicts, contention, and interference. To do so, many approaches have used graph coloring algorithms [11], [25], weighted graph coloring (e.g., to consider traffic load or loss rates) [18], [20], simulated annealing [6], [28], integer programs [4], [12], and even genetic algorithms [7].

3) To evaluate each of the potential configurations without actually reconfiguring all of the networks (impractical), a *predictive quality metric* is typically used. This metric provides the algorithm with a predicted outcome of when assigning the networks a particular set of frequencies. Examples of such metrics include estimating airtime for each network given fairshare and the residual [3], predicting sustained interference [18]–[20], and resulting throughput given the traffic loads [28].

B. *Current Practice and Limitations*

As we briefly discussed in the introduction, many environmental models, spectrum assignment algorithms, and predictive metrics fail to meet the heterogeneous requirements of the spectrum today. We highlight three key shortcomings.

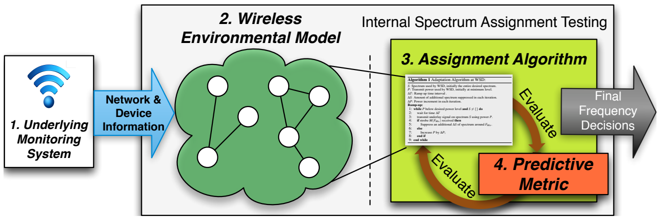

Fig. 1. Common structure of spectrum management systems.

First, a significant amount of prior work in spectrum assignment is *homogeneous* and predominantly 802.11-based. For example, work by Rozner et al. [28], Akl et al. [2], and Murty et al. [22], [23] all assume homogeneous properties while assigning spectrum. That is, they assume overlapping networks will at least coordinate, and that they can all be assigned in the same manner (i.e., they have the same possible set of channels). Even the recent channel changes in 802.11n and 802.11ac would likely require significant changes to these solutions to support the newer standards. However, some of this work does consider different traffic loads (i.e., application-layer requirements) [22], [23], [28].

Second, more recent work that considers heterogeneous technologies is Wi-Fi centric, i.e., the goal is reconfiguring and optimizing a Wi-Fi network to avoid interference from networks using other technologies (e.g., [26]). It does not comprehensively consider or predict interference between all possible heterogeneous radios in the environment. Additionally, their work also does not provide a concrete algorithm on how to reorganize the spectrum (even the 802.11 network) to avoid the heterogeneous interference estimated.

Third, work that comprehensively considers heterogeneous networks (i.e., not Wi-Fi centric) continues to make *overly simplified assumptions* about networks and the RF environment. For example, Peng et al. [25] and Sooyeul et al. [12] make similar critical assumptions untrue of environments with heterogeneous technologies. Both assume that conflicts all have the same weight, i.e., interference from one device is just as severe as from another device (i.e., it is a binary conflict graph). Clearly, this is not true in practice. For example, cordless phones have been shown to reduce 802.11's throughput to near zero [9], whereas ZigBee networks have a lesser impact (e.g., around 60-70% [16]). Given the density of the spectrum, binary conflicts will likely not lend well to efficient configurations.

Additionally, this mentioned work and others (e.g., [7], [11]) incorrectly assume that all radios use the same channels, i.e., center frequencies and bandwidths. This is a critical assumption that is not true in heterogeneous environments: they have different center frequencies, different bandwidths, and even different spectrum band capabilities (e.g., 2.4GHz vs. 5GHz). We believe that this assumption is made to allow the problem to more easily be reduced to a variant of graph coloring (one of the most popular assignment algorithms). We find that without these simplifying assumptions (e.g., of unified channels), heterogeneous environments will likely be more difficult to reduce to basic variants of graph coloring.

More importantly, we found graph coloring to be overly restrictive when trying to formulate and model an environment and spectrum with heterogeneous networks. The most basic form of graph coloring considers conflicts to be binary in weight. Weighted graph coloring has been used to reflect different amounts of interference from partially overlapping channels (e.g., [18]), but we found it difficult to try and capture various degrees of back-off, interference, and asymmetry in a single metric. There are various estimates of coordination and interference needed to properly reconfiguring the spectrum.

III. REQUIREMENTS AND PRINCIPLES OF DESIGN

With a better understanding of spectrum assignment and the limitations in current practice, we present the design requirements needed to support general heterogeneity and the trends of diversity in today's spectrum, followed by our principles of design to satisfy these requirements.

A. Design Requirements To Support Spectrum Trends

To support current trends in the spectrum, the spectrum management system must **support general heterogeneity** between networks and devices, and it must **accommodate evolution** of the protocols and bands over time. The design should *not* be Wi-Fi centric, or make overly simplifying assumptions about the bands, protocols, or channels.

Supporting general heterogeneity between networks and devices is the key to properly organizing the spectrum today. This means that the components in the spectrum management system must represent and account for aspects of diversity across the PHY, MAC, and application-layer:

- PHY Layer: Diversity in terms of the potential bands supported by each radio's physical layer, their center frequencies and bandwidths, and the propagation characteristics based on different transmission powers (i.e., the components must support asymmetric spatial properties).
- MAC Layer: Different access schemes used to coordinate the spectrum must be supported, and importantly: it must be possible to represent each pair of radios in the environment and whether they coordinate based on their MAC properties. Asmmetry is also important, as different MAC layers and/or settings can lead to one radio coordinating with another, but not visa-versa.
- Application Layer: Like prior (albeit homogeneous) work has shown, it is important to consider application layer properties when organizing the spectrum, e.g., traffic load and desired throughput or airtime [22], [23].

Accommodating evolution ensures that, as the protocols, standards, and spectrum bands change over time, the spectrum management system does not require significant changes. The system should be able to support new protocols and bands without major changes to the model, algorithm, or metric.

B. Our Principle of Design and Approach

These requirements highlight the many diverse properties our components must support to accommodate heterogeneity and evolution. Our basic principle of design to meet these requirements is to *describe, represent, and organize the environment* ***using fundamental properties*** *of the spectrum and protocols,* ***remaining protocol independent*** *where possible.*

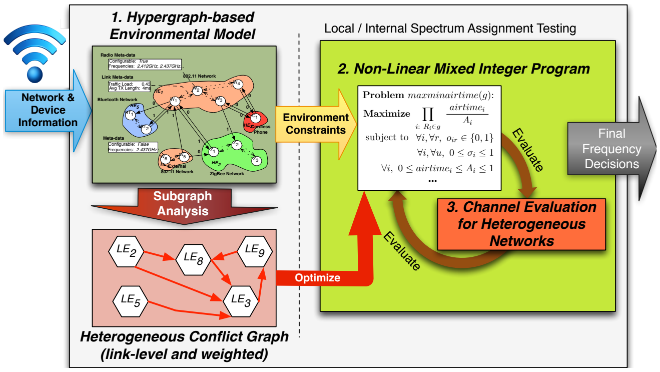

Fig. 2. An overview of our spectrum management system design.

Using fundamental properties: To accommodate heterogeneity and evolution, we design each component of the spectrum management system based on fundamental properties of the spectrum and its protocols, *not* specifics of protocols, standards, or spectrum bands. This ensures that many different protocols and bands can be described using the design, and that new protocols can be supported. For example, our components make no assumptions about "channels" which are specific to standards and technologies. Simplifying assumptions about channels lead to some of the problems in prior work. Instead, our work breaks this PHY-layer property (i.e., channels) in to two fundamentals: 1) A set of center frequencies each radio supports, and 2) An associated bandwidth with these frequencies. This removes specifics of protocols and spectrum bands, while maintaing the system's support of heterogeneity and evolution. If two radios support different spectrum bands, this is represented by their possible set of center frequencies.

Likewise, our system does not describe the many specific details of the PHY and MACs. For example, the details of modulations or whether an 802.11n network is operating in greenfield mode or not (i.e., a specific to a standard). Instead, as another example of a fundamental property, it breaks these properties down in to the fundamentals that matter towards spectrum management: does network X coordinate with network Y given its properties at the PHY and MAC?

Remaining protocol independent *where possible*: While we design the majority of our components to only use fundamental properties in support of heterogeneity and evolution, we believe that there are some particular areas of components where being overly generic sacrifices accuracy or efficiency in assignment. In these areas, we believe that specifics should be used to improve accuracy. If the specifics are not available: the system should provide a generic and reasonable estimate.

IV. HETEROGENEOUS SPECTRUM MANAGEMENT DESIGN

In this section, we present our novel design to accommodate heterogeneity and evolution in the spectrum, meeting the requirements we have discussed in Section III-A, and follows our principles given in Section III-B. This design is illustrated in Figure 2. First, the system takes the information provided by

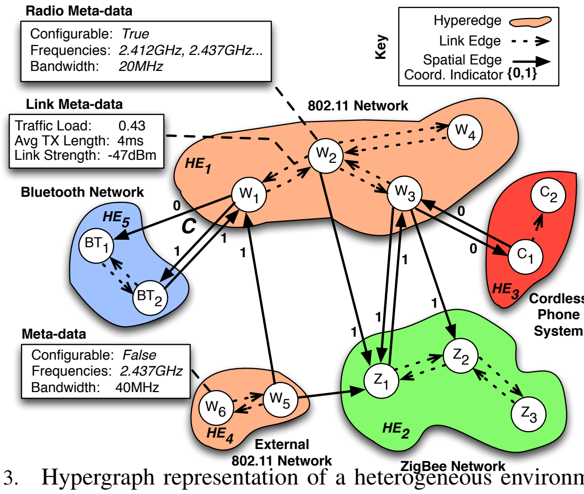

Fig. 3. Hypergraph representation of a heterogeneous environment.

the underlying monitoring system and constructs a *hypergraph-based* environmental model. This model unifies information from the underlying monitor, represents diverse properties of networks, and provides a structure that is easily searchable for conflicts. We introduce subgraph analysis on the hypergraph to detect heterogeneous conflicts, and a novel optimization to mitigate conflicts and interference.

A. *Heterogeneous RF Environmental Model*

Here, we present a highly descriptive hypergraph-based model that represent the key fundamentals in heterogeneous environments. Hypergraphs are a generalization of a graph in which a *hyperedge* (which we abbreviate *HE*) can connect any number of vertices. The purpose and benefit of the hypergraph-based RF environmental model is three fold: 1) To have a structured representation of the RF environment that unifies the information from diverse underlying monitoring infrastructures, guiding them to collect and structure the necessary information such that our system can organize their environment, 2) To have a rich structure that supports various types of constraints in the RF environment occupied by heterogeneous networks, and 3) To structure the RF environment and the behavior within it in a way that is easily searchable for behavior between heterogeneous networks that match conflict types (to optimize based on). We briefly describe the components in our hypergraph, referring to Figure 3 for discussion.

Hypergraph Components & Representation

Vertices: At the base of our hypergraph-model is a set of vertices that represent a wireless *radio*. In todays environments, we believe that it is important to make the "base unit" a *radio* rather than network (or device) for several reasons. First, networks can span larger areas in which different radios receive different levels of interference. A level lower, devices can have multiple heterogeneous radios (e.g., a laptop with a Bluetooth and Wifi radio). This makes devices too coarse-grained.

Edges: Our model has 3 edge types: hyperedges, link edges, and spatial edges. A hyperedge represents a network dependency between radios. For example, the network dependency between W_5 and W_6, represented by hyperedge HE_4. This provides constraints to the algorithm, ensuring that radios within the same network have uniformly chosen frequencies. A link edge represents one-way communication between two radios in our model, denoted $LE\{X,Y\}$ from radio X to radio Y. These edges imply spatial overlap, and link edges can only exist between radios that are connected by a hyperedge (i.e., communication happens within a network only). A spatial edge explicitly models a radio Y being within range of a radio X, denoted $SE\{X,Y\}$. This edge is also uni-directional, an important characteristic as we discussed in our requirements to not assume symmetry. Spatial edges also have a binary indicator that indicate whether the radio Y being within range of radio X causes it to back off from X's transmissions, meta-data also states whether this back-off is digital or power sensed.

Meta-data: Finally, radios, links edges, and spatial edges all carry forms of meta-data. Radios carry meta-data that provides additional constraints in assignment. For example, its possible frequencies, bandwidth, and whether the radio is configurable. Not all radios within range will be configurable, like a neighbor's radios that are within range of a user's home. Links carry meta-data that is used by the assignment algorithm and predictive channel metric. This metadata helps estimate fairshare of airtimes between links that coordinate, and predict heterogeneous interference between links that do not. Spatial edges that denote radio Y is within range of radio X carry meta-data that includes the signal strength of X received at Y.

The hypergraph and its associated meta-data can represent many different types of conflicts in the environment, and asymmetry. We will talk about many of these in the next subsection when we discuss how to derive a conflict graph of heterogeneous conflicts from the hypergraph.

B. *Deriving a Heterogeneous Conflict Graph*

The benefit of having a graph-based model to represent the environment and the interactions within it, is the ability to flexibly search the graph for various specific relationships between components. In our case, to search for various relationships between radios and links indicative of conflicts. This allows us to derive a traditional conflict graph of heterogeneous conflicts from the hypergraph, used as a basis for spectrum assignment.

Building a conflict graph using subgraph isomorphism

To flexibly derive a more traditional conflict graph from the environment of heterogeneous networks represented by our hypergraph, we leverage subgraph isomorphism (also known as subgraph matching), used in many other fields such as social networking, data mining, and anomaly detection to search for specific relationships between entities [8], [15], [24]. We search our larger hypergraph G for subsets of nodes (communication links) with specific attributes and relationships that match a (conflict) subgraph template H.

In Figure 4, we highlight several conflict subgraphs that belong to the set H. Each represents a type of conflict between pairs of communication links (e.g., LE_1 and LE_2). Note that for link LE_2 to conflict with link LE_1, the spatial edge labeled SE_1 *must* exist. Otherwise, the receiver is not impacted. Note that conflicts due to continuous back-off from analog receivers will manifest in our estimation of airtime due to other transmitters.

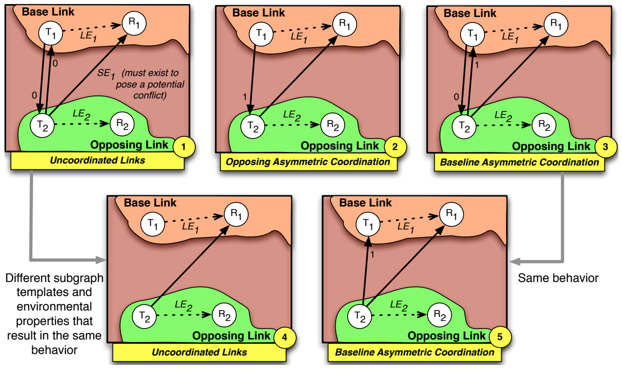

Fig. 4. Examples subgraphs used to detect conflicts in our hypergraph.

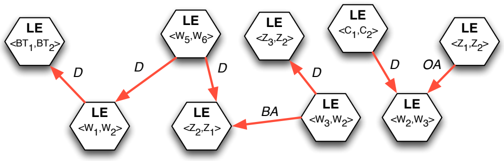

Fig. 5. The resulting conflict graph with conflict annotations.

Each templates varies in the *exact* properties between the two links that represent conflicts. For example, both transmitters being within range but not coordinating (subgraph #1), both transmitters hidden (subgraph #4), and templates for various asymmetric situations (subgraphs #2, #3, #5). Although there are many possible templates needed to match all potential conflicts using subgraph isomorphism, the resulting behavior can be the same. For example, subgraphs #1 and #4 result in the same outcome: both transmitters being unable to coordinate. Subgraphs #3 and #5 both result in the base link coordinating with the opposing link, but not visa-versa.

Knowing the conflict behavior is important in estimating interference, which can be classified in to 3 categories: 1) *Dual uncoordinated links* where both links do not coordinate with each's transmissions, leading to the highest potential number of conflicting transmissions, 2) *Opposing asymmetric coordination* where the interfering link coordinates with the base link, but not visa-versa, 3) *Baseline asymmetric coordination* where the base link coordinates with the opposing link, but not visa-versa.

Using subgraph isomorphism, our subgraph templates, and our categories for annotations, we derive the more traditional conflict graph used as a basis for optimization from our larger hypergraph-based RF environmental model. In Figure 5, we show the resulting conflict graph for our hypergraph example (Figure 3). As expected, vertexes are links, and the edges denote a conflict from an opposing (interfering link) to a base link (with an appropriate annotation). For example, link edge $LE\{W_5,W_6\}$ conflicts with link edge $LE\{Z_2,Z_1\}$, and this conflict results in both transmitters being uncoordinated. We encourage the reader to go through each conflict to gain a better understanding of our templates and the matching process. This graph provides meaningful information about the conflicts in the environment, including their particular scenario to allow us to weight them as a basis for our optimization. Additionally, it is flexible to support other types of conflicts, and it is generic.

C. A Heterogeneous Predictive Channel Quality Metric

The last critical component needed by the assignment algorithm is the predictive channel quality metric. Given a radio and its set of links in a heterogeneous environment, this metric should consider: 1) Contention from coordinating networks and links, and 2) Interference from heterogeneous radios. Consider the performance of a radio R_i and its links on an active frequency f as follows. First, assume that R_i has an average airtime of A_i based on the demand of its links (i.e., those where it is a transmitter), and consider $C_i(f)$ to be the set of radios on frequency f that R_i coordinates with. Given consideration 1, radio R_i will receive an airtime that is at most the maximum of: a) the residual airtime given the radios in $C_i(f)$, and b) its expected fair share with the radios in $C_i(f)$. This first part is a fairly simple estimation, made by many prior works (e.g., [3], [5], [26]). Given consideration 2, this airtime is degraded by active conflicts when operating on frequency f. Denoting $\sigma_i(f) \in [0,1]$ to be the estimated fraction of airtime lost due to these conflicts, the total estimated performance (or "good airtime") of radio R_i on frequency f would be:

$$airtime_i(f) = max(1 - \sum_{c:\, R_c \in C_i(f)} A_c, \frac{1}{|C_i(f)|+1})(1-\sigma_i(f))$$

The key (and non-trivial) challenge is estimating $\sigma_i(f)$, dependent on coordination on each of the radio's links with all other links in range (MAC layer), the traffic loads between each of these competing links (application layer), as well as various PHY layer properties such as the SINR on the links and their modulations which often provide different robustness properties based on their bitrate and error correction.

The basis of our estimation for $\sigma_i(f)$

To estimate $\sigma_i(f)$, we leverage the following key observations and techniques. First, the value of $\sigma_i(f)$ is driven by links that conflict with R_i's links, available in our conflict graph. Then, the severity of these conflicts and their contribution to a higher loss rate in $\sigma_i(f)$ is dependent on whether the conflict is "active" (i.e., are the radios of the links sharing a channel?), how often transmissions on the conflicting links overlap, and how often overlaps cause a transmission loss (e.g., due to SINR properties). Therefore, the probability of loss due to a conflicting link u on a link j that belongs to R_i is:

$$ProbOfLoss_{ju} = ActiveConflict_{ju} * POverlap_{ju} * OLoss_{ju}$$

By estimating the probability loss from each conflicting link with j, we can estimate the total loss rate on a link as:

$$LinkLossRate_j = 1 - \prod_{u:\, M_u \in U_j} 1 - ProbOfLoss_{ju}$$

As a result, the total sustained interference (i.e., loss rate) on R_i that we have denoted $\sigma_i(f)$ would be the sum of the loss rates on each link that belongs to R_i, scaled by that links contribution to the radio's total airtime:

$$\sigma_i(f) = \sum_{j:\, L_j \in K_i} (LinkAirtime_j \;/\; A_i) * LinkLossRate_j$$

Estimating the $POverlap_{ju}$ and $OLoss_{ju}$ components

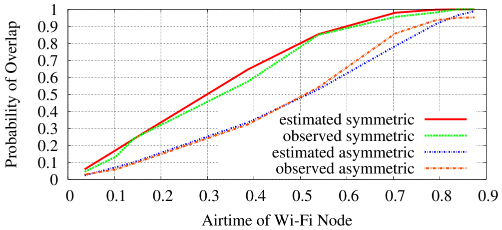

Fig. 6. Evaluating estimated overlaps given (a)symmetric backoff.

Since our assignment system controls what conflicts are active through frequency assignment (i.e., $ActiveConflict_{ju}$), the challenge of estimating $\sigma_i(f)$ is determining the probability of loss due to $POverlap_{ju}$ and $OLoss_{ju}$. Unlike prior work, however, our goal is to estimate these across many technologies and in a way that is generic where possible, and flexible.

To estimate $POverlap_{ju}$, we leverage the observation that without coordination these links operate *entirely independently*. As a result, their events (i.e., packet transmissions) occur continuously and independent of one another. Therefore, if the monitoring system is unable to provide an observed value of $POverlap_{ju}$, we can provide an approximation by modeling the links as independent Poisson processes, using knowledge of their average transmission lengths and airtimes, i.e., fundamentals and not protocol specifics. Our derivation is similar to historical estimations of collision overlap in Ethernet and ALOHA networks without CSMA, however, a key difference is the potential for asymmetry.

First, consider λ_u to be the rate of transmissions from a conflicting link u based on its airtime. Then, consider V_{ju} to be the vulnerability window of transmissions from link j given a conflict with link u. The value of V_{ju} is based on the coordination behavior between j and u, as annotated in our conflict graph. If the annotation on the conflict edge from u to j is D, then both links are uncoordinated meaning u can transmit in the middle of j's transmission, and visa-versa. Therefore, the vulnerability window is $V_{ju} = T_j + T_u$ where T_j is the average transmission time on link j. If j coordinates with u but not visa-versa, the annotation on the conflict edge is BA and the vulnerability window will be $V_{ju} = T_j$, i.e., j is only vulnerable during its transmissions since u cannot sense them. The opposite scenario with an annotation of OA would lead to a vulnerability window of $V_{ju} = T_u$. Given our assumptions, the probability of overlap between j and u would be:

$$POverlap_{ju} = 1 - e^{(-\lambda_u * V_{ju})}$$

To show the importance of modeling asymmetry and that this method can provide a reasonable estimate, we operate Wi-Fi and ZigBee links over the air, controlling their airtimes and coordination behavior. We estimate $POverlap$ and compare it to an actual observed value, shown in Figure 7. Our estimations are close to the observed values. In some cases, the asymmetric overlap rate is half that of the dual uncoordinated ("symmetric") rate. Therefore, if one ignores asymmetric behavior in the environment, the estimate can be significantly more inaccurate.

To accurately estimate $OLoss(l,u)$, we *must* consider specifics of the technologies used by the links since loss during overlap is a factor of many PHY layer properties. This information is not difficult to measure and can be pre-measured and provided by a lookup table. The key considerations to whether an overlap causes a loss between two links is: a) The technologies and their modulations in use, and b) The SINR at the receiver of the base link j. That is, the strength of its reception and the strength of the interference. To illustrate this, we refer the reader to Figure 7 which shows for various reception strengths on a ZigBee link and various interference strengths at the ZigBee link's receiver, what the probability of loss during overlap is with an 802.11 interferer and analog interferer (e.g., cordless phone). Clearly, the outcome depends on the reception and interference strength.

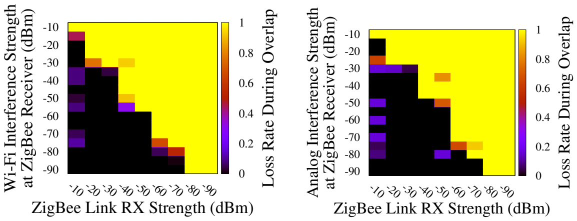

Fig. 7. Heatmap for overlapping transmissions causing a loss.

Using the reception and interference strengths between links in the environment, we can compute lookup tables of overlap loss rates given pairs of technologies and use the values provided by our monitor to index the tables. While this requires N^2 tables, many studies compute these tables to study interference (e.g., [16], [26], [27]). Computing these tables is not complex, yet can ensure accuracy.

D. Spectrum Assignment Algorithm

The last component in our heterogeneous spectrum management system is our assignment algorithm. We introduce a mixed integer program (MIP) that uses the hypergraph-based environmental model (Section IV-A) as input for the constraints in the system (e.g., the potential frequencies of each radio), and our conflict graph (Section IV-B) and metric (Section IV-C) as a basis to optimize. The optimization that we introduce follows our principles of design.

Notation and modeling of our problem

Our representation of the hypergraph and its components are as follows. Let G be the hypergraph of our wireless environment with a total of I radios denoted R_i $(i = 1,...,I)$. Each hyperedge H_e $(e = 1,...,E)$ in graph G contains a set of wireless radios that it connects, such that $H_e = \{R_1, R_2, ...\}$. Next, our graph has S spatial edges SE_s $(s = 1,...,S)$ where SE_{ji} denotes R_i being within range of R_j. If R_i coordinates with transmissions from R_j, then $R_j \in C_i$.

Each communication link edge LE_l (l = 1, ..., L) represents the communication between two radios, where the link edge LE_{ij} denotes a communication link from R_i to R_j. Each communication link $LE_l \in G$ has an average desired airtime of $LinkAirtime_l$ and an average transmission length of T_l. The communication links where R_i is the transmitter belong to the set K_i, and the average airtime A_i of the radio R_i is based on the demand from its links: $A_i = \sum_{j:L_j \in K_i} LinkAirtime_j$. The spectral bandwidth of R_i is denoted B_i, and the possible set of frequencies for each $R_i \in G$ belong to the set F_i. Let the variable f_{ix} represent whether the chosen frequency for R_i is at the index x in F_i (x = 1, ..., X).

$O_{i_x j_z} \in \{0,1\}$ is, for all possible frequency and bandwidth combinations, whether the frequencies F_{ix} and F_{jz} overlap, given respective bandwidths of B_i and B_j. This is used to compute the variable o_{ij}, whether R_i and R_j actively overlap.

Each communication link $LE_l \in G$ has a set U_l of (at least partially) uncoordinated and conflicting links. These links are those that have a conflict edge to LE_l. Since these sets of links are pre-known to the optimization (derived during subgraph analysis), we use the annotations on each link to precompute the vulnerability window V_{ld} for a conflict from L_l to L_d.

Finally, we introduce the variable $airtime_i$ to be our estimated performance of R_i's configuration, its links, and its interactions with potentially active conflicts in the environment. This accounts for airtime sharing and sustained interference (σ_i). We do not include the derivation of σ_i since we have already included the equations and details in Section IV-C. However, we note that o replaces *ActiveConflict* (i.e., a conflict is only active if two radios overlap in the spectrum).

Objective: Our goal is to choose a frequency f_i for each radio $R_i \in G$ to maximize the performance of their networks. In particular, we look to maximize the fraction of airtime received to the airtime desired: $airtime_i/A_i$ for each radio $R_i \in G$. This will not favor certain networks based on desired airtime.

Problem $\max\min airtime(g)$:

$$\textbf{Maximize} \prod_{i:\ R_i \in g} \frac{airtime_i}{A_i}\ , \quad \text{subject to}$$

$$\forall i, airtime_i = min(A_i, max(Residual_i, FairShare_i))(1-\sigma_i) \tag{1}$$

$$\forall i,\ Residual_i = 1 \ - \sum_{c:\ R_c \in C_i} A_c * o_{ic} \tag{2}$$

$$\forall i,\ FairShare_i = 1 \ / \sum_{c:\ R_c \in C_i} o_{ic} \tag{3}$$

$$\forall e : H_e \in g,\ \ \forall R_i \in H_e,\ \ \forall R_j \in H_e,\ \ f_i == f_j \tag{4}$$

$$\forall i, \forall c,\ o_{ic} = \sum_{x=1}^{X_i} \sum_{z=1}^{X_j} O_{i_x c_z} \wedge f_{ix} \wedge f_{cz}\ \in \{0,1\} \tag{5}$$

$$\forall i,\ \sum_{k=1}^{K_i} f_{ik} = 1 \tag{6}$$

$$\forall i, \forall r,\ o_{ir} \in \{0,1\} \tag{7}$$

$$\forall i, \forall u,\ 0 \le \sigma_i \le 1 \tag{8}$$

$$\forall i,\ 0 \le airtime_i \le A_i \le 1 \tag{9}$$

$$\forall i,\ 0 \le Residual_i \le 1 \tag{10}$$

$$\forall i,\ 0 \le FairShare_i \le 1 \tag{11}$$

The program and its constraints: Above, we define a non-linear mixed integer program (MIP) formulation that takes as a parameter g, the hypergraph representation of the environment. We briefly explain each of the constraints and relaxations within. (1) is the constraint on each radio R_i's estimated airtime. The minimum of A_i ensures the radio is never assigned more than its demand. (2) calculates the residual airtime based on all other radios that R_i coordinates with (in C_i). (3) accounts for the fair share of airtime expected from these same coordinating radios. Since o_{ic} is binary, we can count the number of coordinating radios on the same frequency by summing o_{ic} across each radio $R_c \in C_i$. (4) introduces the constraint that radios connected by a hyperedge must have the same chosen frequency. (6) ensures that each radio can only have one center frequency (only one center frequency indicator should be active). The remaining constraints (8 - 11) ensure that many of our variables take on values between 0 and 1. (9) constrains $airtime_i$ to be less than or equal to its demand A_i.

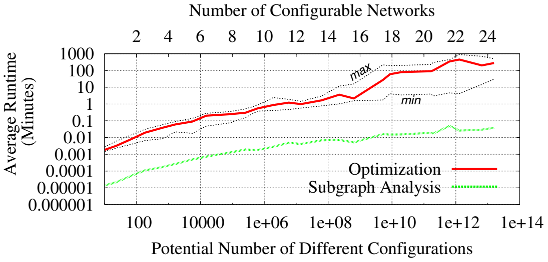

Fig. 8. Runtime of the optimization and subgraph analysis.

V. EVALUATION

Now that we have presented our entire system design, we present an evaluation of our heterogeneous spectrum assignment system. Our goal is to characterize the system's ability to find reasonable and efficient spectrum assignments in heterogeneous environments, improving the performance and fairness of networks within it. We begin with an evaluation of our optimization's runtime, and then present an evaluation of our system in a real heterogeneous testbed.

A. *Performance of the optimization*

An immediate concern is the runtime of our optimization to provide a set of spectrum assignments given various sizes of the environment. In Figure 8, we show the average, minimum, and maximum runtime in minutes given this potential number of configurations, broken down by subgraph analysis time and total runtime. On the top of figure, we also provide an approximate number of configurable networks that result in the potential number of configurations given on the x-axis. As shown, the algorithm will provide what it believes to be an optimal assignment in less than a minute with approximately 1 million different configurations, driven by 10 configurable networks. This is reasonable for the average home environment, small apartment complexes, and possibly larger environments if partitioned in to different workloads. Due to brevity we do not illustrate these results, but we find minimal increase in runtime due to additional static networks. With 40 static networks within range added to the optimization, runtime only increases by approximately 2 minutes.

Recent work has made it possible to parallelize the optimization in terms of memory and computation, with the authors showing reasonable reductions in solve time (e.g., in the cloud) [29]. We have not experimented with this framework as it was released after we completed our evaluation.

B. *Heterogeneous Testbed Evaluation*

Evaluation Methodology: We create several live environments in a heterogeneous testbed where networks are con-

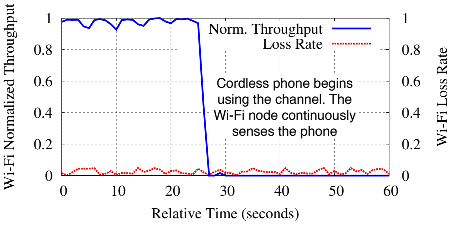

Fig. 9. Cordless phone causing throughput drop on Wi-Fi network due to continual back-off.

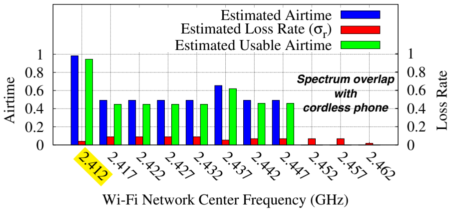

Fig. 10. Estimations with the cordless phone in the upper part of the spectrum.

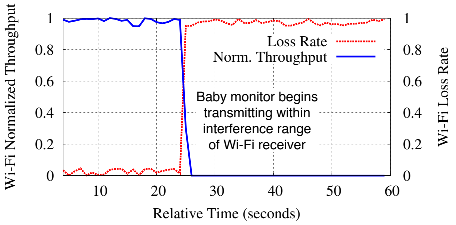

Fig. 11. Baby monitor causing throughput drop on Wi-Fi network due to increased interference.

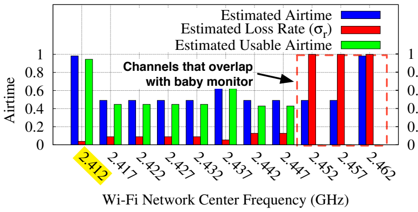

Fig. 12. Estimated airtime, loss, usable airtime for the Wi-Fi network with the baby monitor.

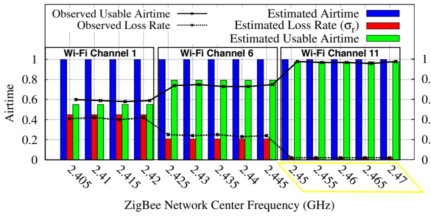

Fig. 13. ZigBee estimations and observed values that properly reflect loss and airtime.

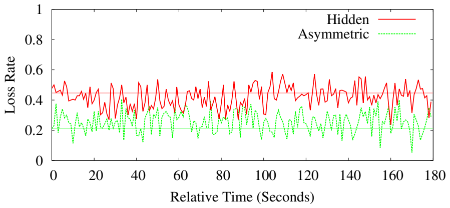

Fig. 14. Estimated and observed loss rates given fully hidden and asymmetric networks.

figured to generate specific demands with various types of conflicts. We classify these expiments in to 4 main categories:

1) *Targeted* scenarios with various types of conflicts between heterogeneous networks that have been reported as common in prior works (e.g., in [16], [26], [27]). These scenarios allow us to provide the reader with a deeper understanding of how our predictive channel metric estimates various types of interference, and how our single algorithm provides assignments that avoid various conflicts.

2) *Weakly constrained* scenarios include multiple heterogeneous networks with the potential for conflicts, however, there is sufficient spectrum to isolate incompatible networks. These scenarios illustrate the spectrum management algorithm's capabilities in relatively isolated environments.

3) *Moderately constrained* scenarios include the configuration of multiple heterogeneous networks with the potential for conflicts present if not placed intelligently, due to a moderately constrained spectrum in terms of demand.

4) *Severely constrained* scenarios are scenarios where frequency isolation is not always possible to avoid all conflicts. Networks must be placed intelligently to reduce interference, reduce contention, and provide fairness across the networks.

Targeted scenarios and evaluation

Starvation and high loss due to *analog transmitters*: We begin with two scenarios that include analog transmitters, often causing starvation in CSMA-networks due to continual back-off and high amounts of loss (observed in many studies [9], [10], [27]). First, we introduce an analog cordless phone within spatial range of a Wi-Fi transmitter. In Figure 9, we plot the normalized throughput and loss rate of the Wi-Fi link before and after we turn the cordless phone on at a relative time of 25 seconds. As shown, the throughput of the Wi-Fi link drops significantly, but the loss rate remains low and stable.

In Figure 10, we break down the channel estimations made by our optimization when deciding how to configure the Wi-Fi network. The *"estimated usable airtime"* is the value of our calculated channel quality metric. However, we also include its two main components: the estimated available airtime due to contention (graphed as "Estimated Airtime"), and the estimated loss rate due to heterogeneous conflicts – our σ value in the main equation (graphed as "Estimated Loss Rate"). As shown, our metric predicts the minimal usable airtime as starvation, *not* loss. Given other active networks, the optimization suggests a center frequency of 2.412 GHz.

Next, we replace the cordless phone with a baby monitor in the spectrum, in to interference range of the Wi-Fi link's receiver (and not transmitter). In Figure 11, we show the result of turning the baby monitor on at a relative time of 25 seconds. Again, the throughput drops significantly, however, loss rate increases in this conflict scenario. As shown in Figure 12, our optimization accurately estimates the result of this conflict scenario: low airtime due to loss, not continual back-off.

Inability to coordinate vs. asymmetric coordination: As discussed, prior studies have reported many asymmetric coordination scenarios between heterogeneous wireless networks, e.g., between devices that have different transmission powers such as ZigBee and Wi-Fi [16], [32]. We evaluate our systems ability to detect and accurately estimate these conflicts.

First, we place a Wi-Fi network on Wi-Fi channel 1 that conflicts with a ZigBee network in our environment where neither backs off to each other. We place another Wi-Fi network on channel 6 that has the same demand as the first Wi-Fi network, however, the distance it is placed at causes the ZigBee network to back-off to the Wi-Fi network, but not visa-versa [16]. We setup another Wi-Fi network on channel 11 with the same demand, but outside of interference range.

We configure the ZigBee network to operate on each of its potential center frequencies, and record the observed usable airtime and loss rates. The resulting channel estimations made by our optimization, as well as the true observed values, are illustrated in Figure 13. As shown, the optimization properly estimates the different loss rates from the Wi-Fi networks, accounting for the specific conflict scenario. The optimization suggests an interference free channel within Wi-Fi channel 11.

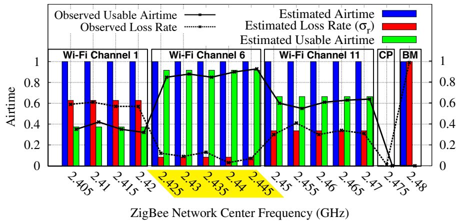

Fig. 15. Optimization accurately assigns a channel given multiple sources of interference.

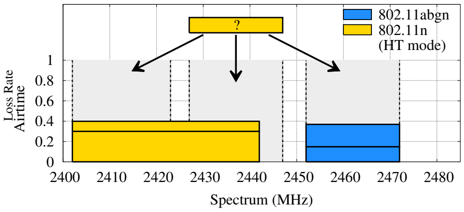

Fig. 16. Digital networks must be carefully assigned and aligned properly.

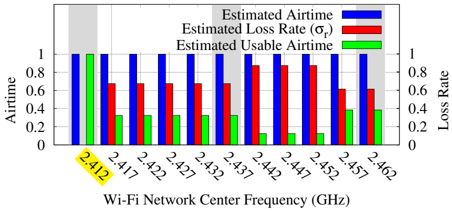

Fig. 17. Our optimization ensures digitally coordinating networks are aligned to coordinate.

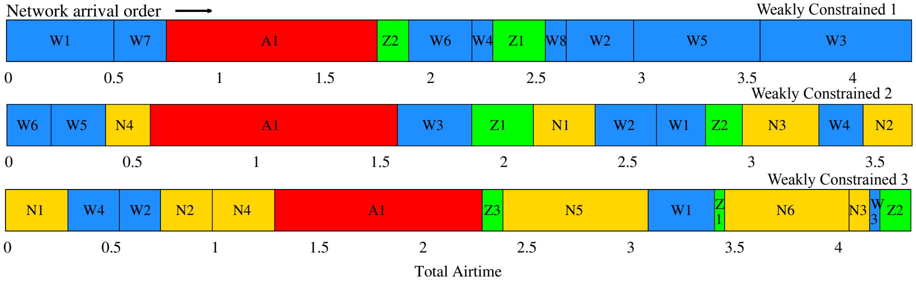

Fig. 18. Weakly constrained scenarios, ordered by arrival for FCFS purposes.

Of course, actual loss rate to fluctuates. We refer the reader to Figure 14 where we plot the estimated loss rates of the hidden and asymmetric scenarios as horizontal flat lines, and the actual observed loss rate over 3 minutes. As shown, the loss fluctuates causing underestimates of the loss of the asymmetric scenario, and slightly overestimate the loss of the hidden scenario. However, they are reasonable estimates and importantly reflect the different conflict scenario.

Accounting for multiple interferers: Next, we setup a targeted scenario to configure a ZigBee network where there are multiple Wi-Fi interferers, and we move the cordless phone and baby monitor to be active in to ZigBee's upper channels. In particular, we setup high sources of Wi-Fi interference on channel 1, moderate sources of interference on channel 11, and low sources of interference on channel 6. In Figure 15, we show that, again, our system provide reasonable approximations of airtimes and loss rates across all of the channels.

Avoiding interference given analog vs. digital coordination: As we described earlier, our optimization and predictive channel quality metric must account for the coordination type when assigning spectrum. We construct the network setup pictured in Figure 16 with two static (unconfigurable) 802.11n 40 MHz networks in HT mode (i.e., they digitally coordinate with a non-legacy preamble), a primary channel of 1, and two static 802.11 networks in legacy mode at channel 11. Then, we introduce a configurable 20 MHz 802.11n HT mode network constrained to the 2.4 GHz band that is within spatial and interference range of the static networks.

Focusing on the non-overlapping channels 1, 6, and 11, our optimization must assign the 20 MHz 802.11n network by properly estimating the impact of it operating with heterogeneous networks and not properly aligning with networks that it digitally coordinates with. We darken non-overlapping channels 1, 6, and 11 in Figures 16 and Figures 17. As shown, if the network is placed in to channel 11, it will receive interference from the legacy 802.11 networks. If placed on channel 6, it will receive interference from the 40 MHz 802.11n networks since it will partially overlap with the 40 MHz network's secondary channel (where it cannot digitally coordinate). If properly aligned with the digitally coordinating 40 MHz networks on channel 1, our optimization accurately estimates no loss to receive its desired airtime.

Weakly constrained scenarios and evaluation

Next we shift our focus to environments that involve configuring multiple networks. Importantly, we compare the organization that our optimization provides against "first-come, first served" organizations that reflect today's environments. That is, networks greedily choose what channel is best for them when they are introduced in the environment, and they are unlikely to reconfigure their assignment (even if a different channel will provide better performance at a later time). When deploying in a FCFS manner, we have the networks greedily choose the channel with the least usage. Additionally, the FCFS assignment is done agnostic to heterogeneous conflicts.

Weakly Constrained 1: We start with a configuration between standard 802.11 networks, two ZigBee networks, and an analog cordless phone that have the respective airtimes shown in Figure 18. In Figure 19, we show the resulting spectrum organization provided by our optimization: heterogeneous networks are isolated, and the 802.11 networks are packed in a way that all networks are able to meet their desired airtime. Alternatively, we show the resulting configuration given a FCFS configuration in Figure 20 (arriving in the order shown in Figure 18). As shown, without careful assignment, heterogeneous networks can share a channel and have performance degraded due to interference. Additionally, homogeneous networks can also receive lower airtime due to non-optimal placement and resulting contention.

Weakly Constrained 2: Next, we reconfigure some of the legacy supporting 802.11 networks to be high-throughput 802.11n networks i.e., they use the newer digital preamble that makes them incompatible with legacy networks. The

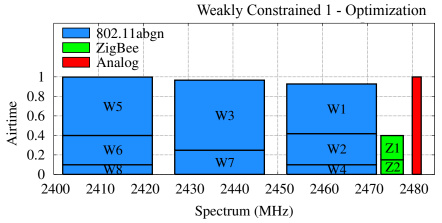

Fig. 19. Our optimization isolates heterogeneous technologies when possible.

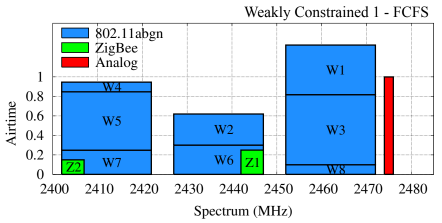

Fig. 20. FCFS configuration where networks greedily choose their channel assignments.

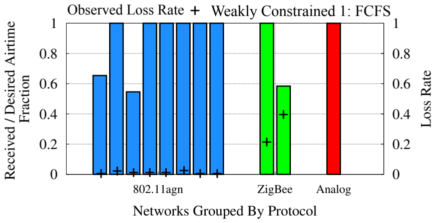

Fig. 21. FCFS observed loss rates and airtimes, normalized by desired airtime.

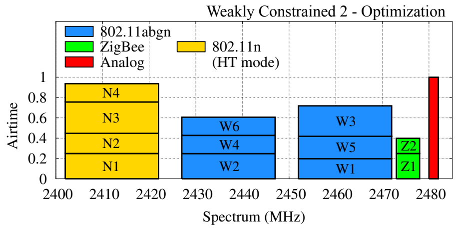

Fig. 22. Heterogeneous technologies are again separated with enough spectrum available.

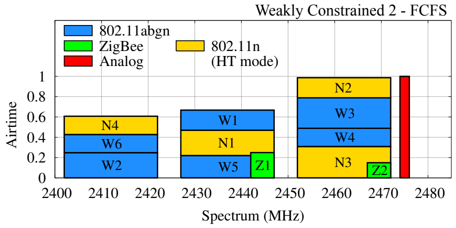

Fig. 23. Resulting assignments given FCFS in the weakly constrained scenario 2.

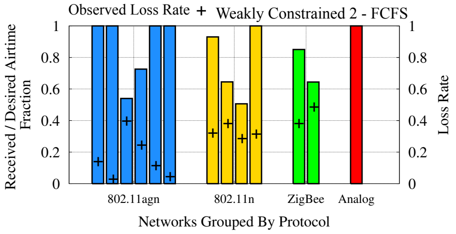

Fig. 24. Higher loss rates and lower received airtime's with FCFS, as expected.

resulting set of networks and their airtimes are shown in Figure 18. Leveraging the resulting conflict graph from our management system that accurately reflects conflicts between the digitally coordinating networks and legacy networks, our optimization continues to isolate the conflicting networks as shown in Figure 22. Comparing to a FCFS assignment shown in Figure 23, again, the inability of the heterogeneous networks to sense each other leads to a mixed spectrum of conflicting networks, that leads to the observed high loss rates and reduced airtimes shown in Figure 24. The legacy supporting 802.11 networks receive slightly lower loss rates due to their use of basic spectrum sensing that allows them to back-off to the high-throughput networks. However, the high-throughput networks receive higher loss rates due to their strict digital-only sensing. The ZigBee networks also receive high loss rates from both sets of 802.11 networks. Although one of these ZigBee networks is still able to meet its desired airtime, losses lead to retransmissions, which can affect the ZigBee network's performance in terms of power consumption.

Weakly Constrained 3: In our final weakly constrained scenario, we introduce the networks in our testbed shown in Figure 18: 40 MHz high-throughput Wi-Fi networks, and configure some of these networks to be able to use the 5 GHz band. The resulting configuration given our assignment system is shown in Figure 25: heterogeneous networks are again isolated, as well as digitally coordinating networks being properly aligned i.e., the 20 MHz high-throughput 802.11 networks are aligned with the primary channel of the 40 MHz networks. Our optimization avoids placing the 20 MHz high-throughput networks in the secondary channel of the 40 MHz networks, which would lead to their inability to coordinate. Due to potential contention, the networks that support the 5 GHz frequencies are pushed in to the 5 GHz spectrum band.

The FCFS configuration is shown in Figure 26, with its observed airtimes and loss rates in Figure 27. Again, the FCFS configuration suffers from low airtime fractions, and high loss rates. In particular, the 40 MHz 802.11n networks suffer greatly due to interference from legacy networks, as well as interference from the 802.11n HT network that resides in the secondary channels. Although the first ZigBee network receives its desired airtime due to retransmissions, its loss rate is over 50% where retransmissions affect power consumption.

Moderately constrained scenarios and evaluation

In our moderately constrained scenarios, we increase the spectrum usage such that demand requires intelligent solutions where simple isolation is not possible. Due to the inefficiencies in FCFS configurations shown extensively in our results, we focus solely on our optimization's results in these scenarios. We will revisit FCFS configurations (as well as largest-first insertion) in the severely constrained scenarios.

Moderately Constrained 1: In the first moderately constrained scenario, we reconfigure Weakly Constrained Scenario 1 such that the analog cordless phone operates in the ZigBee network's higher potential channels, forcing the ZigBee networks to operate within the 802.11 channels to avoid complete starvation from overlapping with the cordless phone. We show the resulting configuration in Figure 28 and ask the reader to compare the original assignments of the 802.11 networks with those in Figure 19. Our optimization reconfigures its assignment of the 802.11 networks to allow ZigBee to operate in its lower channels interference-free. We draw conflict arrows between the 802.11 networks and the ZigBee networks to highlight this. Networks $W5$, $W7$, and $W8$ were reconfigured to share a channel, creating a channel where $Z1$ would receive no interference from conflicts. Likewise, $W3$ and $W2$ were moved together. Additionally, our optimization still considered contention between the Wi-Fi networks and placed them in a way that still allowed them to meet their desired airtimes.

Moderately Constrained 2: Next, we take Weakly Constrained Scenario 3 and add two additional legacy 802.11 networks to it, pushing the demand higher. If one places these two networks in 802.11 channel 11 to operate with the other legacy networks shown in Figure 25, all legacy networks will

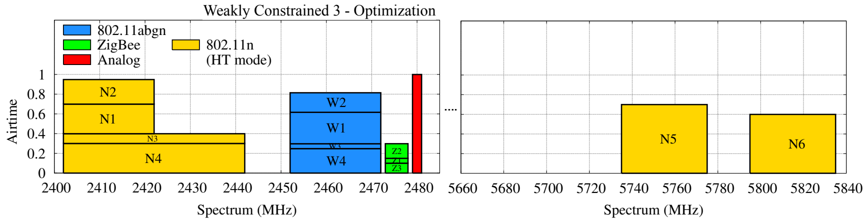

Fig. 25. Network placement given our optimization. The scenario includes 40 MHz networks and the use of high-throughput 802.11 networks in the 5 GHz spectrum band.

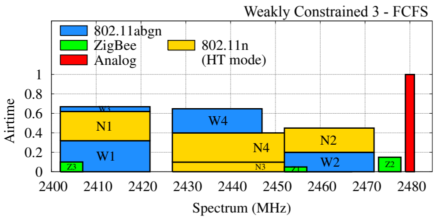

Fig. 26. Significant overlap of heterogeneous technologies in FCFS configuration.

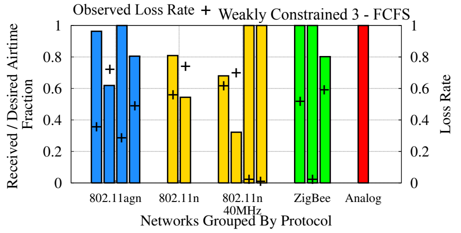

Fig. 27. High loss rates and low performance shown in FCFS configuration.

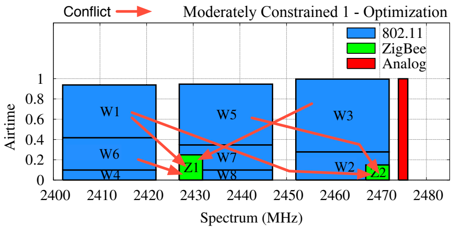

Fig. 28. Our optimization reconfigures 802.11 networks to avoid ZigBee interference.

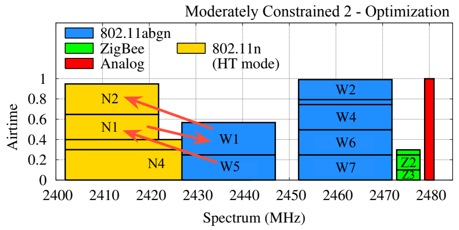

Fig. 29. Optimization intelligently reassigns spectrum to intelligently avoid conflict also.

receive a reduced airtime due to contention. Our optimization determines that 802.11 networks $W1$ and $W5$ (where $W5$ was newly introduced), do not conflict with the 40 MHz networks allowing them to operate in their secondary channel freely. We show the resulting configuration in Figure 29 and draw the conflicts to show the intelligent placement which avoids conflict. Networks $W1$ and $W5$ received no observed loss rates, as well as the two 40 MHz networks not receiving loss from these networks. This shows that under more constrained spectrum, our optimization can still find efficient organizations.

Severely constrained scenarios and evaluation

Finally, we introduce severely constrained scenarios where networks must be intelligently placed to avoid as many conflicts as possible. We compare to FCFS configurations, and two new points of comparison: 1) A largest-first insertion of networks in to the spectrum that uses our predictive channel quality metric, and 2) Our optimization with a Jain's fairness-based objective function. The prior comparison allows us to compare to intelligently placing networks without the runtime of the full optimization, and the latter comparison allows us to compare to an objective of fairness across the networks.

Severely Constrained 1: We begin with the configuration of networks shown in Figure 30, and the resulting airtimes of the 4 different assignment methods in Figure 31. The figure shows, for each method, the fraction of networks that receive at least a certain amount of airtime. First, our optimization with the standard objective finds a configuration that has the most networks meet their desired airtime: 75% of networks. Comparing to the FCFS configuration where only 43% of networks meet their desired airtime. Even with our predictive channel quality metric, a largest-first method of insertion and assignment cannot meet the performance of our optimization. Note that the organization provided by our optimization with the standard objective results in a few networks that receive relatively lower (un-fair) airtime fractions. Using the Jain's fairness objective with our optimization improves the least performing network, but to achieve this several better performing networks receive lesser airtime fractions: only 58% of networks receive their desired demand with the Jain's fairness objective, compared to the 75% with the standard objective.

In Figure 32, we show the resulting loss rates for the networks, focusing on the optimization with the standard objective and the FCFS method of assignment. As shown, neither method is able to avoid complete loss across all of the networks, however our optimization ensures that no network received a loss rate greater than 8%. The FCFS method of assignment has some networks receive significant loss rates between 80-90%. This shows the benefit of our assignment system, as well as its flexibility of different objectives..

Severely Constrained 2: We conclude our evaluation results with the severely constrained scenario of networks depicted in Figure 30. There is a high degree of heterogeneity and overall desired airtime. We show the resulting airtime fractions of the networks given our four insertion methods in Figure 33. Again, our optimization with the standard objective finds a configuration where the most networks are able to meet their desired airtime, where 82% of networks meet their demand. A FCFS configuration has less than half of the networks meeting their demand, and 60% of networks meet their demand using our optimization with the Jain's fairness objective and the largest first method of insertion. The Jain's fairness objective continues to provide the best worst-case performance.

In Figure 34, we compare the loss rates observed using our optimization with the standard objective as compared to the FCFS insertion. With severely constrained networks given their demand and the available spectrum, as well as the high degree of heterogeneity, FCFS is shown to result in high loss rates for many of the networks. Only 20% of networks receive a loss rate lower than 40%, with some networks receiving loss rates greater than 80% again. This is compared to our optimization where high loss rates are still observed (e.g., greater than 50%), but 80% of networks still receive minimal loss rates.

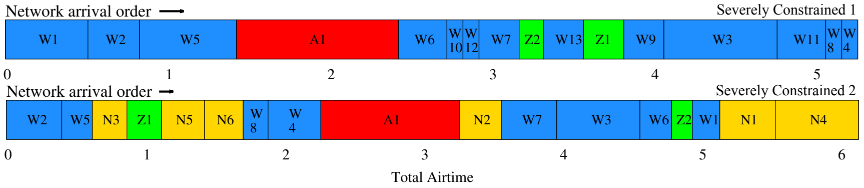

Fig. 30. Networks in the severely constrained scenarios, ordered by arrival for FCFS.

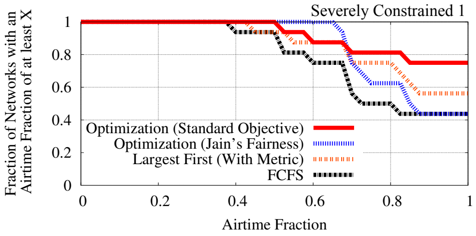

Fig. 31. On average, our optimization with default objective outperforms all other techniques.

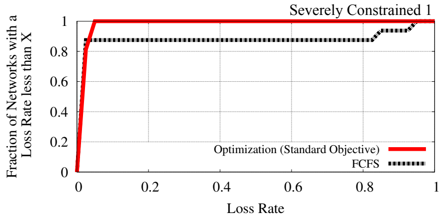

Fig. 32. Our optimization avoid significantly higher loss rates that are observed in FCFS.

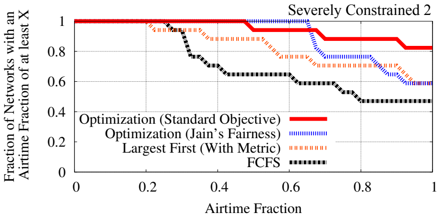

Fig. 33. With a higher degree of heterogeneity, our optimization further outperforms.

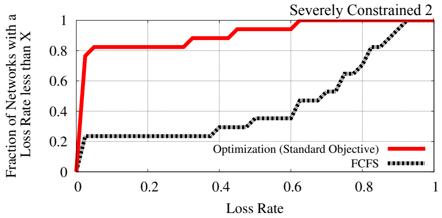

Fig. 34. Our optimization continues to avoid higher loss rates.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a spectrum management system to organize environments with heterogeneous networks. Our system is designed to support various technologies and their evolution over time by describing fundamentals and remaining generic where possible. Through evaluation in a real heterogeneous testbed with various layouts, we showed our optimization can avoid common conflicts (with our single algorithm), and reduce interference in more dense environments. Given the runtime of our algorithm, we believe that it is suitable for home environments and smaller complexes.

As future work, there are additional aspects of the spectrum that can be integrated in to our system. For example, we assume that the signal characteristics (e.g., propagation) are the same on all frequencies. We do not model a drop in signal quality when re-assigning to a higher frequency (e.g., from 2.4GHz to 5GHz), or an improvement in quality when re-assigning to a lower frequency (e.g., from 2.4GHz to the white space TV bands). Pur MIP-based approach provides enough flexibility to add these constraints in the future.

REFERENCES

[1] I. 802.19. Workgroup on standards for coexistence between wireless standards of unlicensed devices. http://www.ieee802.org/19/.
[2] R. Akl and A. Arepally. Dynamic channel assignment in ieee 802.11 networks. In *Portable Information Devices, 2007. PORTABLE07.*
[3] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. SIGCOMM '09.
[4] M. Cesana, D. Maniezzo, P. Bergamo, and M. Gerla. Interference aware (ia) mac: an enhancement to ieee802.11b dcf. In *VTC '05*.
[5] Y.-C. Cheng et al. Jigsaw: solving the puzzle of enterprise 802.11 analysis. SIGCOMM '06.
[6] D.-R. Din et al. A Simulated Annealing Algorithm for Extended Cell Assignment Problem in a Wireless ATM Network. In *Applications of Evolutionary Computing Lecture Notes in Computer Science*, 2001.
[7] Y. Ding, Y. Huang, G. Zeng, and L. Xiao. Channel assignment with partially overlapping channels in wireless mesh networks. WICON '08.
[8] L. Getoor and C. P. Diehl. Link mining: a survey. SIGKDD '05.
[9] S. Gollakota et al. Clearing the rf smog: making 802.11n robust to cross-technology interference. SIGCOMM '11.
[10] R. Gummadi et al. Understanding and mitigating the impact of rf interference on 802.11 networks. SIGCOMM '07.
[11] M. Hassan and A. Chickadel. A Review of Interference Reduction in Wireless Networks Using Graph Coloring Methods. GRAPH-HOC '11.
[12] S. Im and H. Lee. Dynamic spectrum allocation based on binary integer programming under interference graph. In *PIMRC 2012*.
[13] K. Jain et al. Impact of interference on multi-hop wireless network performance. MobiCom '03.
[14] A. Kumar. Enhancing coexistence, quality of service, and energy performance in dynamic spectrum access networks. In *MobiCom*, 2004.
[15] J. Leskovec, A. Singh, and J. Kleinberg. Patterns of influence in a recommendation network. PAKDD'06.
[16] C.-J. M. Liang, N. B. Priyantha, J. Liu, and A. Terzis. Surviving wi-fi interference in low power zigbee networks. SenSys '10.
[17] X. Liu et al. Dirc: Increasing indoor wireless capacity using directional antennas. In *Proceedings of the ACM SIGCOMM*, Aug. 2009.
[18] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted coloring based channel assignment for wlans. *SIGMOBILE Comput. Commun. Rev.*
[19] A. Mishra et al. Exploiting partially overlapping channels in wireless networks: turning a peril into an advantage. IMC '05.
[20] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh. Partially-overlapped Channels not considered harmful. In *ACM Sigmetrics*, 2006.
[21] S. Mishra et al. Detect and avoid: an ultra-wideband/wimax coexistence mechanism. *Communications Magazine, IEEE*, 2007.
[22] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing high performance enterprise wi-fi networks. NSDI'08.
[23] R. Murty, J. Padhye, A. Wolman, and M. Welsh. Dyson: an architecture for extensible wireless lans. USENIXATC'10.
[24] C. C. Noble et al. Graph-based anomaly detection. KDD '03.
[25] C. Peng et al. Utilization and fairness in spectrum assignment for opportunistic spectrum access. *Mob. Netw. Appl.*, 11(4):555–576, 2006.
[26] S. Rayanchu et al. Catching whales and minnows using wifinet: Deconstructing non-wifi interference using wifi hardware. NSDI'12.
[27] S. Rayanchu, A. Patro, and S. Banerjee. Airshark: detecting non-wifi rf devices using commodity wifi hardware. IMC '11.
[28] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. Traffic-aware channel assignment in enterprise wireless lans. In *ICNP 2007*.
[29] Y. Shinano et al. FiberSCIP - A shared memory parallelization of SCIP. ZIB-Report 13-55, September 2013.
[30] V. Shrivastava et al. CENTAUR: realizing the full potential of centralized wlans through a hybrid data path. In *ACM MobiCom*, 2009.
[31] M. Vutukuru, K. Jamieson, and H. Balakrishnan. Harnessing exposed terminals in wireless networks. In *In USENIX NSDI*, 2008.
[32] X. Zhang and K. G. Shin. Enabling coexistence of heterogeneous wireless systems: case for zigbee and wifi. MobiHoc '11.