

VillageLink: Wide-Area Wireless Coverage

Veljko Pejovic*, David L. Johnson[†], Mariya Zheleva[‡], Elizabeth M. Belding[‡] and Albert Lysko[†]

*School of Computer Science, University of Birmingham, UK

Email: v.pejovic@cs.bham.ac.uk

[†]Council for Scientific and Industrial Research, South Africa

Email: {djohnson, alysko}@csir.co.za

[‡]University of California, Santa Barbara, CA, USA

Email: {mariya, ebelding}@cs.ucsb.edu

Abstract—White spaces promise to revolutionize the way wireless connectivity is delivered over wide areas. However, large-scale white space networks face the problem of allocating channels to multiple contending users in the wide white space band. To tackle the issue, we first examine wireless propagation in a long-distance outdoor white space testbed and find that a complex combination of free-space loss and antenna effects impacts transmission in white spaces. Thus, a need arises for a strategy that goes beyond simple channel utilization balancing, and uses frequency probing to profile channels according to their propagation properties. We devise VillageLink, a Gibbs sampling-based method that optimizes channel allocation in a distributed manner with a minimum number of channel switching events. Through extensive simulations we demonstrate that VillageLink results in a significant capacity improvement over alternative solutions.

I. INTRODUCTION

Internet connectivity is available to merely 39% of the world’s population [1]. The main cause of limited Internet penetration stems from the fact that more than three billion people live in rural areas. These areas are hard to connect via copper cables, fiber optic or cell phone base stations due to high deployment cost and low population density which renders these techniques economically infeasible. Rural areas are also hard to reach via cheap license-free solutions such as WiFi, as these technologies, operating in 2.4 or 5 GHz bands, have a very limited connectivity range.

In the 50-800 MHz band, a large block of frequencies has recently been freed due to the analog to digital TV transition. This spectrum, called white spaces, promises to deliver an affordable means of providing wide area coverage. It is extremely attractive for rural areas as the propagation range is an order of magnitude higher than in the bands used by competing technologies. However, a distributed, resource-efficient solution for network organization, especially for spectrum allocation within a network, is needed for further proliferation of rural area white space deployments.

The goal of channel allocation is to assign one of the channels to each of the network nodes. Traditionally, the issue has been expressed as the graph coloring problem where a color (channel) is assigned to a node so that network interference is minimized, and consequently the capacity is maximized. In a network operating over a small set of available frequencies, such as WiFi, channels generally do not exhibit significant differences in terms of propagation properties. White space

networks, however, operate over a very large span of frequencies, and propagation properties can vary drastically over these channels. Channel assignment in such a network has to satisfy conflicting goals: maximize useful transmission by preferring channels with superior propagation, and minimize interference by favoring channels that propagate over a shorter radius.

Since we propose white space connectivity for impoverished regions, we concentrate on making our solution as cost efficient as possible. Therefore, we propose to reuse the existing TV antennas already installed in even the most remote rural areas. Unfortunately, this further complicates the problem of channel allocation as these antennas exhibit uneven and unpredictable propagation behavior over the wide white space spectrum. Any analytical solution that provides a clear picture of frequency quality becomes impossible, and a direct inference of propagation properties is needed.

In this paper we successfully address the above challenges by designing a light-weight frequency profiling methodology to evaluate channel quality and a novel channel allocation method that assigns operating frequencies to base stations with the goal of minimizing the impact of interference over the useful signal levels in a network. We compile these contributions into a practical channel profiling and allocation scheme for wide area white space networks called VillageLink. We test VillageLink’s frequency probing mechanism on a long-distance software-defined radio white space link we deployed and confirm that antenna effects and the environment are a significant reason for high propagation diversity among white space channels. Through simulations we evaluate VillageLink’s channel allocation. We show that our frequency-aware channel allocation leads to up to twice as much network capacity than an alternative heuristic based on interference avoidance, and that with its high performance, efficient resource usage and distributed nature, represents a practical solution for wide area white space coverage in rural areas.

II. WIDE-AREA WHITE SPACE NETWORKS

White spaces represent a historic opportunity to revolutionize wide area wireless networking. White spaces not only deliver much greater communication range than Gigahertz frequencies, they also support non-line of sight communication, including transmission through vegetation and small obstacles, which makes them highly suitable for various terrain




Fig. 1. Analysis of received signal strength over the UHF band. The received signal strength is difficult to predict as it is dependent on a complex mix of antenna gain patterns, cable issues and environmental structures around antennas and between the transmitter and receiver.

configurations. However, white space networks have to deal with unique peculiarities of transmission over a wide band of relatively low frequencies, and should enable license-free unplanned deployments in rural developing regions.

A. Wide band frequency selectivity

The variation in the free-space loss across a band is termed “dynamic range” and is calculated as $D_{dB} = 20 \log(f_U/f_L)$ where f_L and f_U are the lowest and the highest frequency in the band, respectively. White space frequencies operate on a wide band of low frequencies, leading to a large dynamic range. Thus, unlike in traditional wireless networks, such as WiFi or GSM, free-space loss over white spaces is not uniform over the range of frequencies on which a network operates. Besides wide dynamic range, white space links often experience uneven fading due to antenna patterns and the environment. The fractional bandwidth (FB) for a frequency band, calculated as a ratio of operating bandwidth and the central frequency, determines how wideband an antenna should be in order to have the same gain over all frequencies with the band. Again, due to the low central frequency and a large operating band, white spaces require significantly wider band antennas than GSM and WiFi. Such antennas are hard to design; high gain across the full frequency range is nearly impossible. Consequently, white space links are prone to the effects of imperfect antennas and surrounding structures.

To confirm this, we deployed a 3 km outdoor non line-of-sight link using USRP2 radios, standard UHF log periodic TV antennas and a 1 W amplifier at the transmitter. The transmitter was configured to send a series of 10 consecutive 1MHz wide BPSK modulated PN sequence probes across the full UHF white space spectrum while avoiding existing TV bands. The receiver was synchronized to the transmitter clock and scanned the spectrum both with the transmitter turned off (baseline scan) and with the transmitter sending probes (signal scan).

Figure 1 shows the received signal strength across the UHF TV band. Three TV stations were detected and probes did not occur at these frequencies. The received signal strength does not fall off monotonically with increasing frequency, which would be the case if only free-space loss determined the propagation loss. Instead, due to the antenna characteristics, cabling and the environment, the propagation loss is non-uniform




Fig. 2. Simulated antenna gain for the antenna used in the outdoor link.

across the UHF band. Using the WIPL-D antenna modelling software we created a model of the deployed antenna. The results are shown in Figure 2. While an antenna with no surrounding structures has a more predictable gain pattern, when surrounding structures and antenna imperfections, such as bent or missing elements, are introduced, the antenna gain pattern has far less predictability. Residential TV Antenna installations often require long lengths of low-grade coaxial cable which can easily get damaged during installation by home users. Analysis of periodicity of dips in the received signal strength in Figure 1 indicates a cable that has been crushed or bent tightly approximately 2 m from the connector. Predicting the type of TV antenna being used, the structures surrounding the antenna or cable imperfections is not possible and provides strong support for frequency probing in white spaces.

In addition to antenna effects, a part of frequency selectivity may stem from the environment and terrain. *Shadowing*, slow fading due to physical obstacles on the signal path, would still be detected and accounted for with frequency probing. Unlike shadowing, *multipath*, which leads to rapid variation of propagation within a channel, cannot be captured through our current probing method. However, channel allocation only requires knowledge of the average channel gain of the channel that is captured by frequency probing.

B. Channel assignment in white space networks

The problem of channel assignment in wireless networks is often expressed with graph coloring, where each color represents a different channel. For a link, one of the available central frequencies is assigned so that a goal, such as maximum throughput, is achieved. In the channel allocation literature on traditional wireless networks all colors are considered equal in terms of their propagation properties [13]. However, in white spaces, the transmission range varies significantly among frequencies in the band due to the wide dynamic range and antenna effects (Figure 2). Therefore, *selection of the operating frequency can impact the existence of a link itself*. This further complicates the problem of graph coloring, as now not all colors are equal. The choice of the color effects the graph structure, thus the existing approaches to frequency assignment are not directly applicable.

C. Network Architecture

The network scenario that describes the setting in which VillageLink will operate is given in Figure 3. In this paper,




Fig. 3. Layout of a targeted white space network showing interference scenarios between TV and white spaces, and between white space networks in different domains. White space BSs within the same domain send BS to BS probes (BBPs) to calculate the channel conditions among themselves.

we consider wide-area white space networks that consist of individual base stations (BSs), each with a set of associated customer-premise equipment (CPE) clients. We term one such BS with its CPEs a *cell*. A BS and all the CPEs within a cell operate on the same channel; thus, when considering channel allocation we use “BS” and “cell” interchangeably. All cells that operate within the same administration are called a WRAN domain. The existence of TV transmission and other interfering white space networks not in our service set reduces the number of channels available to the BSs within our domain. The aim of our work is to develop a channel allocation algorithm so that the overall network performance within our WRAN domain is maximized.

III. CHANNEL PROBING AND MEDIUM ACCESS FOR WIDE-AREA NETWORKS

Channel probing is an important tool for propagation evaluation over a wide white space frequency range. Unfortunately, the existing MAC protocols proposed for wide area networks [12], [11], [3] do not explicitly support frequency profiling. The MAC protocol that most closely resembles our proposed system is IEEE 802.22. The protocol has built-in protection for primary users and mechanisms to move to new channels but has no built-in mechanism to choose from a set of available channels. Instead of rebuilding an entire MAC layer, we propose to extend the 802.22 protocol to include a feature that performs frequency profiling on all available channels. In our frequency profiling scheme we measure the SNR of the probe that was received at a BS using the previously measured power spectral density of channel with no probes and the power spectral density measured when a probe is present. We assume that the network nodes are static.

In order to perform frequency profiling between BSs and between BS and CPEs on all available channels, a mechanism is required to coordinate probing timing for channel probe senders and listeners. We use a token mechanism, in which a BS only sends channel probes when it has a token. All other BSs and CPEs without tokens are in a listening state. When a BS is in a probing state, it sequentially steps through the full white space TV channel set and only sends a probe on non-interfering channels.

The 802.22 specification makes use of CPEs to sense for primary users and extend the sensing coverage area. We propose to use a similar notion when listening for probes. CPEs of one BS experience interference from all other BSs. To account for this CPEs can be instructed to listen for probes from BSs with whom they are not associated. Frequency profiling results for CPEs are sent back to the associated BS on the final upstream frame once the CPE has listened on the full set of white space channels. The average SNR value of a received probe heard at a BS and its associated CPEs is used to incorporate the average interference on the system. These SNR values, from each of the cells that received the probes, are unicast on the back channel to the sending BS. Results received at sending BSs are distributed to neighbouring BSs, where two BSs are defined as neighboring if a probe can be exchanged between them on at least one frequency.

Once the probing process is completed each BS stores a “localized set” of information on signal propagation at different frequencies: 1) within its own cell, obtained through aggregation of probing results from the cell’s CPEs; 2) between itself and other BSs and associated CPEs that overheard probes; and 3) within cells that are served by neighboring BSs.

A. Calculating probe SNR

At each channel c_i the probe SNR with no TV interference can be calculated using the average power in the probe listen window and the noise level from an initial scan when no probes are present. From the measured SNR we can extract the channel gain:

$$H(c_i) = \frac{SNR(c_i) \cdot N_0 \cdot W}{Pt} \quad (1)$$

where N_0 , W and Pt denote the noise constant, channel width and the transmission power, respectively.

These SNR measurements are able to capture BS to CPE channel gains and channel gains (interference) between a BS and other cells that include the BS and associated CPEs. We do not capture CPE to CPE interference (client uplink effecting another client downlink on the same channel) as this would require CPEs to carry out probes and not allow our system to scale. Moreover, BS superframes are time-synchronized in 802.22 and only a small portion of the uplink frame is likely to overlap with the downlink frame of another cell. In addition, 802.22 is also able to move into a co-existence mode with an adjacent cell experiencing interference in which frames are fully coordinated between cells. We also assume channel reciprocity for our measurements – a common assumption for systems using the same channel for up and down links.

IV. CHANNEL ALLOCATION

In this section we devise a distributed channel allocation algorithm that uses information obtained through frequency profiling (Section III) and does not incur channel switching overhead typical for other allocation schemes. Our approach is based on the annealed Gibbs sampler [2], a technique that can help us minimize a target function in a distributed way.

A. Gibbs Sampling

The Gibbs sampler is a Markov chain Monte Carlo (MCMC) technique for obtaining random samples from a multivariate probability distribution. The sampler is useful in situations where the joint distribution is unknown or difficult to sample, but the conditional distributions of variables are known and easy to sample. The Gibbs sampler draws samples from a multivariate probability distribution $p(x_1, \dots, x_N)$ as follows:

- Initialize all variables x_1, \dots, x_N to (random) starting values x_1^0, \dots, x_N^0 .
- In every iteration $j = 1..k$, sample each variable x_i from the conditional distribution $p(x_i | x_1^j, \dots, x_{i-1}^j, x_{i+1}^{j-1}, \dots, x_N^{j-1})$ to obtain x_i^j .

After the above process is finished, we are left with $x_1^j, \dots, x_N^j; j \in [1..k]$ samples from the joint distribution p .

We can solve the channel allocation problem through Gibbs sampling, if we obtain the samples from a distribution that: *i*) is related to overall network performance, *ii*) depends on the selected operating channel of each of the BSs, *iii*) isolates the impact of each of the BSs on the total optimization function, *iv*) can be calculated in a distributed way and sampled independently at each of the BSs, *v*) favors states that lead to maximum performance. Next, we develop a network performance metric that can be used as a basis for a distribution that satisfies these demands.

B. Network Performance Metric

Traditionally, the goal of a channel allocation protocol is to assign available channels to BSs so that the total network capacity is maximized. The capacity $C_i(c_p)$ of a single cell operating on the channel c_i is:

$$C_i(c_i) = \sum_{k \in \mathbb{K}_i} W_k \log(1 + SINR_{ik}(c_i))$$

where \mathbb{K}_i is the set of CPEs within the cell, W_k is the width of a part of the channel c_i used by CPE k , and $SINR_{ik}(c_i)$ is the signal to interference plus noise ratio at the CPE k . We approximate the presence of all clients within the cell with a single *virtual CPE* with an SINR value $SINR_i(c_i) = \sum_k SINR_{ik}(c_i) / |\mathbb{K}|$. The cell capacity is now:

$$C_i(c_p) = W \log(1 + SINR_i(c_i))$$

where W is the full channel width, essentially a sum of all W_k as a cell operates in an OFDMA mode. This approximation hides channel distribution within the cell and helps us concentrate on inter-cell interaction.

If we consider a network with N cells, with a given channel assignment $\mathbf{c} = (c_1, c_2, \dots, c_N)$, $c_i \in \mathbb{C}$, where \mathbb{C} is the set of available channels, the total network throughput is a sum of all individual capacities at their respective allocated channels:

$$C(\mathbf{c}) = \sum_i C_i(c_i) = \sum_i W \log(1 + SINR_i(c_i)) \quad (2)$$

A single BS's decision on the operating channel changes the interference level at all its neighboring BSs. In the above

equation the interference is accounted for in the SINR, which is embedded within the logarithmic function. Thus, the impact of a single BS on the total sum is hard to isolate, and the total capacity is not a suitable metric for distributed computation.

One of the ways to circumvent this is to revert to a tighter problem formulation that prevents interfering BSs from concurrent transmission [6]. While this can be enforced in a network that employs carrier sensing and collision avoidance, in our setting long distances between BSs render such coordination inefficient [12]. In addition, allowing some interference often yields more capacity than restricting concurrent transmissions [9]. Another approach is to modify the optimization function and instead of maximizing capacity minimize total interference [5]. However, in white spaces, where available channels can differ drastically in terms of their propagation properties, *a channel allocation that leads to minimal interference may not necessarily lead to maximum capacity*.

We propose a novel network performance metric that allows distributed performance optimization with Gibbs sampling, and term it *cumulative interference plus noise to signal ratio (CINSR)*. It represents a sum of inverse of SINR experienced at each of the cells. CINSR can be seen as the overall ratio of the impact of harmful factors, noise and interference, to the beneficial one, received signal strength. Thus, our goal is to minimize it:

$$CINSR(\mathbf{c}) = \sum_{i=1}^N \frac{1}{SINR_i(c_i)} \quad (3)$$

$$= \sum_{i=1}^N \frac{N_0 W + \sum_{j=1..N, j \neq i} ch(i, j) PH_{ji}(c_i)}{PH_i(c_i)} \quad (4)$$

The first term in the numerator within the above sum is the thermal noise, whereas the second term is the sum of interference experienced at cell i , and originating from all other BSs that transmit at the same channel. Interference from a single source is a product of P - the transmission power and $H_{ji}(c_i)$ - the propagation gain from BS j to cell i on channel c_i . The function $ch(i, j)$ is equal to 1 if i and j operate on the same channel, and otherwise it is equal to 0. The denominator in the above equation is the average signal strength received by the clients of the BS i ; the average channel gain from BS i towards the clients on channel c_i is denoted by $H_i(c_i)$.

We now isolate the impact of a single BS i on $CINSR(\mathbf{c})$ and term it *local CINSR*:

$$CINSR_i(\mathbf{c}) = \frac{N_0 W}{PH_i(c_i)} \quad (5)$$

$$+ \sum_{j \neq i} ch(i, j) \left(\frac{PH_{ij}(c_i)}{PH_i(c_i)} + \frac{PH_{ji}(c_i)}{PH_j(c_i)} \right) \quad (6)$$

Information needed for $CINSR_i(\mathbf{c})$ calculation, namely $PH_i(c_i)$, $PH_{ji}(c_i)$, $PH_{ij}(c_i)$ and $PH_j(c_i)$, is available locally at BS i , through channel probing described in Section III.

C. The Gibbs distribution

The Gibbs distribution associated with the function $CINSR$ and a positive temperature T is the probability distribution on c^N (the combined channel state space of all BSs) defined as:

$$\pi(\mathbf{c}) = \frac{e^{-CINSR(\mathbf{c})/T}}{\sum_{\mathbf{c}' \in c^N} e^{-CINSR(\mathbf{c}')/T}} \quad (7)$$

The above distribution is of special interest as it favors states in which $CINSR$ is low. In addition, the channel selected by BS i is independent of all non-neighboring BSs and the distribution fulfils all the conditions listed Section IV-A.

The Gibbs sampler draws a sequence of samples from the above distribution by having each of the BSs i independently sample its local Gibbs distribution $\pi_i(\mathbf{c})$:

$$\pi_i(\mathbf{c}) = \pi_i(c_i | \mathbf{c}_{\setminus i}) = \frac{e^{-CINSR_i(c_i | \mathbf{c}_{\setminus i})/T}}{\sum_{\mathbf{c}' \in c^N} e^{-CINSR_i(c'_i | \mathbf{c}_{\setminus i})/T}} \quad (8)$$

and transitions to the sampled local state, converging to the stationary distribution $\pi(\mathbf{c})$ (see Section IV-E). Here $\mathbf{c}_{\setminus i}$ denotes a fixed assignment of channels to all BSs but BS i .

Distribution $\pi(\mathbf{c})$ favors low $CINSR$ states when the temperature is low. While our goal is to minimize $CINSR$, by keeping the temperature low we risk getting stuck in a local minimum early in the process. *The annealed Gibbs sampler* introduces a slow decrease of temperature T to zero according to a *cooling schedule*. Therefore, in the beginning the probability of exploring a wide range of states is high, and as the time goes to infinity, the procedure converges to the minimum $CINSR$ state.

D. Channel Allocation Algorithm

Algorithm 1 VillageLink channel allocation – distributed

```

1: {Executed at the base station  $i$ }
2: while  $t < t_{end}$  do
3:    $T = f(T_0, t)$  { $f$  - schedule,  $T_0$  - starting temperature}
4:   for all channel  $c'_i \in \mathbb{C}$  do
5:      $\mathbf{c}' = (c_1, c_2, \dots, c'_i, \dots, c_N)$ 
6:     Calculate  $CINSR_i(\mathbf{c}')$ 
7:   end for
8:   for all channel  $c'_i \in \mathbb{C}$  do
9:      $\mathbf{c}' = (c_1, c_2, \dots, c'_i, \dots, c_N)$ 
10:    Calculate  $\pi_i(\mathbf{c}')$ 
11:   end for
12:   Sample a random variable according to the law  $\pi_i$  and choose the next channel of the BS  $i$  accordingly.
13:   Send information about the newly selected channel to  $i$ 's neighbors.
14: end while
15: Switch the network interface to the last selected channel.

```

Algorithm 1 is executed at each of the BSs. The temperature falls off with time, ensuring that the Gibbs sampler converges towards the global minimum of $CINSR$. The starting time for all the BSs has to be loosely aligned, and can be achieved with a standard synchronization scheme such as NTP.

Compared to some other distributed channel allocation schemes [5], [8], Algorithm 1 has an attractive property that no channel switching is needed until the convergence. To see why note that the calculation of the local $CINSR$ is done after the probing process, and during the algorithm run the only

variable parameter is $ch(i, j)$. At BS i this parameter can be updated irrespective of the actual operating channel of BS j . In every step a BS decides on its current channel and sends the decision to its neighbors, who then update their $ch(i, j)$ tables. Once the cooling schedule is completed BSs switch to their channel of choice (line 15 in Algorithm 1). This greatly speeds up the convergence, as the channel allocation process is not limited by the channel switching time.

E. Algorithm convergence

Convergence of a Gibbs sampler, and its annealed version, is a well researched topic [2]. Here we prove the convergence of our method, indicating that it is a natural heuristic for solving the channel allocation problem.

Proposition 1: The Gibbs distribution π (equation 7) represents a Markov random field.

Proof: A Gibbs potential V associates a real number $V_\Gamma(\mathbf{s})$ with each subset Γ of a set S . The potential is determined by the state \mathbf{s} of the nodes in Γ and is defined as zero if Γ is not a clique. An energy function $\mathcal{E}(\mathbf{s})$ maps each of the graph states to a real number. We say that the energy function *derives from the potential* V if:

$$\mathcal{E}(\mathbf{s}) = \sum_{\Gamma} V_\Gamma(\mathbf{s}) \quad (9)$$

where the summation goes over all subsets of the set S . The Gibbs distribution where the energy derives from a Gibbs potential is a Markov random field (pg. 260 in [2]), and we proceed with showing that the function that we use to construct the Gibbs distribution in equation 7 – $CINSR(c)$ derives from the Gibbs potential.

We can represent $CINSR$ as a sum of local impact of cliques of the graph of BSs \mathcal{A} . $CINSR$ then takes the form described by equation 9 and can be used as the energy function for Gibbs sampling:

$$\begin{aligned} CINSR(\mathbf{c}) &= \sum_{i \in \mathcal{A}} \frac{N_0 W}{PH_i(c_i)} + \\ &+ \sum_{\{i, j\} \in \mathcal{A}} ch(i, j) \left(\frac{PH_{ij}(c_i)}{PH_i(c_i)} + \frac{PH_{ji}(c_i)}{PH_j(c_i)} \right) \\ &= \sum_{\mathcal{B} \subset \mathcal{A}} V_{\mathcal{B}}(\mathbf{c}) \end{aligned}$$

Here V denotes the Gibbs potential. The potential is defined for all subsets \mathcal{B} of the set of BSs \mathcal{A} as:

$$V_{\mathcal{B}}(\mathbf{c}) = \begin{cases} N_0 W / PH_i(c_p) & \text{if } \mathcal{B} = \{i\} \\ ch(i, j) \left(\frac{PH_{ij}(c_i)}{PH_i(c_i)} + \frac{PH_{ji}(c_i)}{PH_j(c_i)} \right) & \text{if } \mathcal{B} = \{i, j\} \\ 0 & \text{if } |\mathcal{B}| \geq 3 \end{cases}$$

Note that the potential is non zero only for cliques of size one and two. Thus, energy $CINSR(\mathbf{c})$ derives from the Gibbs potential and, consequently π is a Markov random field. ■

For a network of N BSs, each running a Gibbs sampler over its local Gibbs distribution $\pi_i(\mathbf{c})$, channel allocation converges

in variation¹ towards the Gibbs distribution π , since the process can be described as a Gibbs sampler on a finite state homogeneous Markov chain represented by the selected channel allocation, for which the Gibbs distribution (equation 7) is the invariant probability measure (Example 6.5, pg. 288 in [2]). Note that direct sampling of the capacity (equation 2) does not provide any guarantees on the performance as the capacity equation cannot be transformed to an energy function that derives from the Gibbs potential. Thus, we develop *CINSR*. Finally, for a fixed network of N BSs implementing Algorithm 1, channel allocation converges in variation towards a limit distribution that only puts positive probability mass on the states of minimum global energy, as we rely on the annealed Gibbs sampler (example 8.8, pg. 311 in [2]). Conditions that the cooling schedule has to satisfy in order for convergence to happen can be found in [4].

V. EVALUATION

The VillageLink system consists of our frequency profiling method built on top of the 802.22 MAC protocol, and the channel allocation algorithm based on Gibbs sampling. Experimental evaluation of such a system is challenging due to the need for a wide area outdoor deployment. In addition, off-the-shelf 802.22 equipment is not yet commercially available, and software defined radio platforms cannot support the synchronization that the MAC protocol requires [10]. Therefore, we evaluate our protocol in a simulated setting. However, the initial experimental investigation of channel probing and frequency selectivity in white spaces, presented in Section II-A, was performed on a 3 km outdoor link.

A. Simulation Setup

For a comprehensive evaluation of the channel allocation algorithm, we rely on a Matlab-based custom simulator. We explicitly take into account high variability of signal propagation in the white space band by modeling propagation with the Friis transmission equation, which also figures antenna gains, transmission power and distance between antennas. Earlier, in our outdoor testbed, we confirmed that frequency dependence of antenna gain is the most dominant factor that leads to the frequency diversity in white spaces (Figure 1), thus we model antenna effects in detail.

We use publicly available antenna models (www.hdtvprimer.com/ANTENNAS/comparing.html) and the Numerical Electromagnetics Code (www.nec2.org) antenna modeling software to examine propagation over different frequencies with different antennas. Figures 4(a) and 4(b) show the radiation patterns seen from the center frequency (598 MHz) of the white space band for two different antennas. In Figure 5 we plot frequency dependence of antenna gain. We found that the shape of the antenna pattern does not change significantly for different frequencies. The gain, on the other hand, changes significantly and unpredictably, as seen in Figure 5. Thus, in the simulations we use the antenna




Fig. 4. Radiation patterns of Winegrad PR9032 UHF Yagi/corner reflector antenna used as a base station antenna, and AntennasDirect DB-2 2-Bay UHF antenna; one of the client antennas used in our evaluation.




Fig. 5. Antenna profiles of four of the antennas used in our evaluation. One of the profiles, Winegrad PR-9032, corresponds to the BS antenna; the other three correspond to client antennas.

pattern shape of the center frequency to account for antenna orientation, and we use the full gain over frequency diversity.

All BSs in our simulations use the Yagi antenna from Figure 4(a), as this antenna exhibits the best performance of all the antennas that were modeled. In our simulation we assume clients make use of existing TV antennas used to receive terrestrial TV broadcast signals. Operators have no control over the variety of antennas used by clients and we randomly select antennas from a set of 17 possible client TV antennas ranging from outdoor Yagi antennas with a gain of 15dBi to simple indoor loop antennas with a gain of 3dBi.

We run our experiments over a white space band from 443 MHz to 875 MHz as the antenna models we use perform reasonably well within this range. The band is divided into 36 TV channels, each 6 MHz wide, with a 6 MHz guard band between adjacent channels. In all the experiments we simulate a $100km \times 100km$ field with random BS placement and random antenna orientation. Each of the BSs has a single associated virtual client at a distance uniformly picked from 0.2 km to 20 km and with its antenna pointed directly towards the BS. We also simulate a TV station that covers a part of the field with its signal and occupies two adjacent channels.

B. Channel Allocation Convergence

Our solution to channel allocation represents a heuristic as we cannot give guarantees on how long will it take for the MCMC process to reach the target invariant distribution. To gauge the practical behavior we simulate *Algorithm 1* in a network of ten BSs and five available channels. We

¹Convergence in variation describes convergence of an array of samples to a probability distribution and is defined in [2], pg. 128.




Fig. 6. Algorithm convergence with the (a) exponential, and (b) logarithmic cooling schedule. Each line corresponds to a different starting temperature.

are interested in the algorithm convergence under different Gibbs sampling parameters. We experiment with two common cooling schedules: *a)* logarithmic: $T = T_0 / \log(t + 2)$ and *b)* exponential $T = T_0 \alpha^t$. Here T denotes the temperature at time t , T_0 is the starting temperature, and $\alpha \in [0, 1]$ is a real number (set to 0.995 in our experiments).

The selection of the starting temperature is important for proper annealing. In Figure 6 we plot total network capacity achieved with the two schedules and four different starting temperatures for each. The impact of the starting temperature is clearly visible: the higher T_0 is, the more time it takes for the algorithm to converge. At the same time, higher temperatures ensure exploration of a large part of the solution space, and generally lead to a better solution. We can also see that $T_0 = 10^{-6}$ does not result in any variation of capacity as the algorithm progresses – the sampler is “frozen” and BSs will stick to the initial channel allocation without exploring the full solution space. There is a trade-off, dictated by T_0 , between the convergence time and the assurance that the optimal value will be found. In the rest of the paper we fix T_0 to 1, a value that allows full exploration of the solution space and converges in a reasonable amount of time, and concentrate on the exponential schedule as it exhibits much faster convergence.

C. CINSR as a Performance Metric

To confirm that *CINSR* is a good choice for the network performance metric, we compare it with an alternative – overall interference and noise in the network – used as a metric in the allocation algorithm proposed by Kauffmann et al. [5]:

$$I(\mathbf{c}) = \sum_{i=1}^N \left(N_0 W + \sum_{j=1..N, j \neq i} ch(i, j) PH_{ij}(c_p) \right) \quad (10)$$

The impact of a single BS on the sum is defined as *the local interference*:

$$I_i(\mathbf{c}) = N_0 W + \sum_{j=1..N, j \neq i} ch(i, j) (PH_{ij}(c_p) + PH_{ji}(c_p)) \quad (11)$$

We modify equation 7 and equation 8 to include $I_i(\mathbf{c})$ instead of $CINSR_i(\mathbf{c})$, and $I_i(\mathbf{c})$ instead of $CINSR_i(\mathbf{c})$, respectively. The necessary conditions for the Gibbs sampler convergence still hold, and we apply an algorithm analogous to Algorithm 1. Note that, defined this way, the interference function still uses the results of channel probing, yet it does




Fig. 7. Comparison of the total network capacity achieved with CINSR and Interference metrics. We simulate under-provisioned and over-provisioned number of channels with respect to the number of base stations in the network.




Fig. 8. Total network capacity with varying number of channels and BSs.

not account for the balance between well propagating channels that are preferred by the CPEs and inferior channels that minimize inter-cell interference.

1) *Channel under-provisioning*: We simulate a network with a number of contending BSs higher than the number of available channels, a typical case in the urban developed world. We put 50 cells in the same $100km \times 100km$ region. We experiment with a varying number of available channels. The total network capacity is plotted in Figure 7(a). When multiple cells operate on the same frequency the network is in a low SINR mode, and capacity can be increased by interference minimization. From Figure 7(a) we see that the two versions of the Gibbs sampler perform equally well with a small number of available channels. As we increase the amount of available spectrum, BSs have more freedom to operate at different channels with minimal interference. Therefore, frequency-dependent performance of CPEs associated with the BSs becomes an important factor that impacts total capacity. Since this factor is not accounted for in equation 10, this version of the Gibbs sampler results in a channel allocation that delivers less capacity than the version that uses *CINSR*.

2) *Channel over-provisioning*: We now fix the number of available channels to 36 and compare the performance of the two versions of the algorithm with the number of BSs varying from 5 to 35. The total network capacity is plotted in Figure 7(b). When the number of channels is greater than the number of BSs there is more than one allocation that leads to minimal interference. However, not all of the allocations are favored by the CPEs. Through the factor $H_i(c_p)$ *CINSR* accounts for the frequency dependent intra-cell preferences, and assigns channels that maximize capacity within each of the cells. The results presented here point out that *channel allocation in white spaces remains important even in rural*




Fig. 9. Fairness with varying number of channels and base stations (the closer the fairness index value is to one, the better).

areas where the channel availability is high.

D. Comparison to alternative channel allocation methods

Channel allocation is a difficult problem to solve in a distributed setting. Heuristics are often used instead of a rigorous solution and we compare our approach with: 1) *Least congested channel search (LCCS)* - a heuristic where each of the BSs individually scans for a channel with the least number of other BSs assigned to it [8]; 2) *Preferred intra-cell channel allocation (PICA)* - in this greedy method each of the BSs selects the channel for which it observes the highest channel gain towards its own CPEs ($\arg \max_{c_p} H_i(c_p)$).

We compare the experimental behavior of different solutions in a number of scenarios encompassing various numbers of BSs and available white space channels. We run each of the algorithms 100 times in each of the scenarios.

1) *Total network capacity*: In Figure 8 we plot the total network capacity as we increase the number of cells in the system from 5 to 50. To ensure consistency among points in the graph, we do not generate a new topology every time we increase the number of cells, but add randomly placed cells to the existing topology. Each of the topology sequences are evaluated in environments with 10, and 20 available channels. We plot average values and two standard deviations (represented by error bars) for each data point. VillageLink performs better or equal to the alternatives in all scenarios. The benefits of frequency-probing based channel allocation grow with the number of cells. In some scenarios, such as 50 BSs - 10 channels, VillageLink delivers twice as much capacity as the next best alternative, LCCS.

2) *Fairness*: In Figure 9 we plot the Jain fairness index for cell capacity with channel allocations determined by VillageLink, LCCS and PICA. We plot average values and two standard deviations (represented by error bars) for each data point. Although we designed VillageLink as a method to optimize total network capacity, it also ensures a remarkably fair allocation of resources. As the number of cells grows, the fairness of VillageLink is more pronounced as it stays close to 1 while the fairness indices of PICA and LCCS drop.

VI. RELATED WORK

Over the past decade most efforts to provide broadband connectivity to remote rural are based on modified WiFi [12]. Propagation in White spaces bands are drastically different to WiFi bands and networking protocols need be reconsidered

for new white space spectrum. Ma and Tsang [6] use an integer linear programming solution for frequency allocation to deal with channel heterogeneity in white spaces. However, frequency reuse is restricted to well defined interference domains where no two BSs are allowed to transmit at the same time. Motivated by [9], we rely on a more sophisticated representation of interference, measuring its impact through probing to account for interference during the allocation process. Gibbs sampling was applied to distributed channel allocation, client association and power control [5], [7]. VillageLink differs from previous approaches by incorporating frequency dependence of both useful signal transmissions and interference. In addition, VillageLink minimizes channel switching and information exchange among nodes.

VII. CONCLUSION

White space networks are largely unexplored, and their straightforward implementation might prove difficult due to unique characteristics they exhibit. In this work we show how the heterogeneity of white space frequencies imposes unique challenges when it comes to channel allocation in a wireless network. Rather than simply minimizing interference, a channel allocation policy has to account for transmission quality over different channels as well. To tackle the problem we develop VillageLink, a channel allocation protocol that relies on the knowledge of signal propagation in the whole white space band before it performs distributed channel assignment that converges towards a network-wide optimum.

REFERENCES

- [1] *ITU World Telecommunication/ICT Indicators Database*, 2013.
- [2] P. Bremaud, editor. *Markov Chains, Gibbs Fields, Monte Carlo Simulation and Queues*. Springer, 1999.
- [3] C. Cordeiro, K. Challapali, D. Birru, B. Manor, and S. Diego. IEEE 802.22: An Introduction to the First Wireless Standard based on Cognitive Radios. *Journal of Communications*, 1(1):38–47, 2006.
- [4] B. Hajek. Cooling Schedules for Optimal Annealing. *Mathematics of Operations Research*, 13:311–319, 1988.
- [5] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks. In *INFOCOM'07*, May 2007.
- [6] M. Ma and D. H. Tsang. Joint Design of Spectrum Sharing and Routing with Channel Heterogeneity in Cognitive Radio Networks. *Physical Communication*, 2(1-2):127–137, March 2009.
- [7] V. Mhatre, K. Papagiannaki, and F. Baccelli. Interference Mitigation through Power Control in High Density 802.11 WLANs. In *INFOCOM'07*, Anchorage, AK, May 2007.
- [8] A. Mishra, S. Banjaree, and W. Arbaugh. Weighted Coloring Based Channel Assignment for WLANs. *Mobile Comput. Commun. Rev.*, 9:19–31, 2005.
- [9] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-Based Models. In *HotNets'06*, 2006.
- [10] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. Steenkiste. Enabling MAC Protocol Implementations on Software-Defined Radios. In *NSDI'09*, Boston, MA, April 2009.
- [11] R. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *NSDI'07*, April 2007.
- [12] B. Raman and K. Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *MobiCom'05*, Cologne, Germany, August/September 2005.
- [13] I. Stojmenovic, editor. *Handbook of Wireless Networks and Mobile Computing*. Wiley-Interscience, 1st edition, February 2002.