

Chapter 9

VillageLink: A Channel Allocation Technique for Wide-Area White Space Networks

Veljko Pejovic, David Lloyd Johnson, Mariya Zheleva, Elizabeth M. Belding and Albert Lysko

This chapter is an extension of V. Pejovic, D. L. Johnson, M. Zheleva, E. M. Belding and A. Lysko, VillageLink: Wide-Area Wireless Coverage, COMSNETS'14, Bangalore, India, January 2014.

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.

V. Pejovic
School of Computer Science, University of Birmingham, Birmingham, UK

D.L. Johnson (✉) · A. Lysko
Meraka Institute, CSIR, Pretoria, South Africa
e-mail: djohnson@csir.co.za

M. Zheleva · E.M. Belding
Department of Computer Science, University of California, Santa Barbara, CA, USA

© Springer International Publishing Switzerland 2015
A.K. Mishra and D.L. Johnson (eds.), *White Space Communication*,
Signals and Communication Technology, DOI 10.1007/978-3-319-08747-4_9

249

9.1 Introduction

Internet connectivity is available to merely 39% of the world's population [7]. Lack of internet access is primarily restricted to developing regions, however, many remote communities in the developed world are currently disconnected as well [22]. 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, the complexity of signal propagation over the wide white space band, and the economic necessity of a resource-efficient, unlicensed, distributed solution for rural areas renders wide area white space networks uniquely challenging to realize.

In a white space network a number of frequency selective effects will be present due to the topology, vegetation and antenna design. Because of the highly complex nature of electromagnetic propagation and the difficulty in accurately modeling fading, the analytical solution that provides a clear picture of frequency quality in white spaces is not practical. Frequency profile, however, is extremely important as it can be used as a basis for channel assignment in a white space network. However, even when frequency propagation information is available, channel allocation in a white space network is very hard. In a network where the span of available channels is not large, such as in WiFi networks, channel allocation can be cast to a graph coloring problem. In white spaces, the wide range of available channels leads to drastic differences in propagation among channels. These differences stem from the variation of free space propagation over frequencies, but also from antenna properties, as in practice antennas do not perform uniformly over a very wide span of white space frequencies. Finally, frequency assignment in such a wide band network has to satisfy conflicting goals: maximize useful transmission by allocating frequencies with superior propagation properties and minimize interference by allocating frequencies that propagate over a shorter radius.

In this chapter we successfully address the above challenges by designing a lightweight 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 envi-

ronment 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, VillageLink represents a practical solution for wide area white space coverage in rural areas.

9.2 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 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.

9.2.1 TV Spectrum Availability

White spaces spectrum can also be used by primary users, typically television and wireless microphones. In order to understand the amount of available white space spectrum, we conducted a number of spectrum scans in rural and urban regions in South Africa, Zambia and the USA.

South Africa has five TV channels utilizing VHF and UHF TV bands; all channels are available in urban centres and a portion of these channels available in rural areas. Zambia only has two national TV programs with very limited coverage. The USA, in contrast, has anywhere between 10 and 25 terrestrial TV channels available depending on location. We carried out spectrum scans in rural and urban areas in the USA and Southern Africa using an Ettus USRP2,¹ 8dBi log-period antenna, and a laptop. More details on the measurement setup can be found in [12]. Figure 9.1 shows the spectrum scan of the lower portion of the UHF band used for television broadcasting. Urban areas of the USA have very limited white space available. Rural areas of the USA are comparable to urban areas of South Africa. Rural areas of Zambia and South Africa have an abundance of white space and are thus very well suited for rural connectivity solutions.

¹ www.ettus.com.

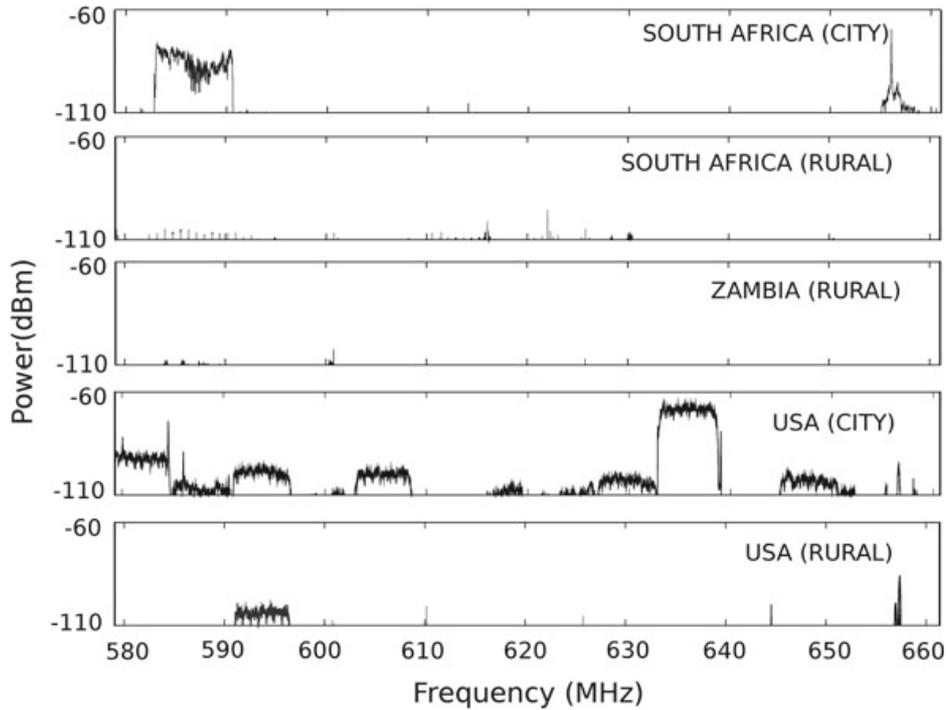


Fig. 9.1 Analysis of available spectrum in the lower UHF band for urban and rural areas in South Africa, Zambia and the USA

9.2.2 Wide Band Frequency Selectivity

One of the unique challenges of establishing wireless communication in the TV bands is the large amount of frequency selectivity due to variation in gain and free space loss across the operating band. This frequency selectivity complicates the choice of the optimal communication channel.

The large variation in free-space loss across the band also known as “dynamic range”. Dynamic range is calculated as follows:

$$D_{dB} = 20 \log (f_U / f_L) \quad (9.1)$$

where f_L and f_U are the lowest and the highest frequency in the band, respectively. In Table 9.1 we summarize the dynamic range of a number of traditional wireless systems. Free-space loss in a traditional wireless network, such as WiFi or GSM, is relatively uniform over the range of frequencies these networks operate on. The reason for low dynamic range in these networks lies in the fact that they either operate over a relatively narrow band of frequencies, such as 50 MHz for GSM and 80 MHz for 2.4 GHz WiFi, or they operate on high central frequencies where the difference between the lowest and the highest frequency diminishes, as is the case with 5 GHz WiFi. White spaces, however, operate on a wide band of low frequencies, and the difference in propagation between white space frequencies can be large. Creating an antenna which has a flat frequency response across the entire white space band is,

Table 9.1 Dynamic range and fractional bandwidth of different wireless systems

Technology	f_L (MHz)	f_U (MHz)	D (dB)	FB (%)
802.11 (2.4GHz)	2,412	2,484	0.26	2.9
802.11 (5GHz)	5,170	5,700	0.85	9.8
GSM900	935	960	0.23	2.6
White spaces	43.25	797.25	25.31	179

**Fig. 9.2** White space 3 km experimental link in sub-urban area of Pretoria, South Africa between CSIR offices and a staff house

however, challenging. Note that the same issue does not arise in GSM (as well as 3G and 4G/LTE) networks, that can also operate on a wide range of frequencies (e.g. GSM850, GSM900, GSM1800). Unlike with white spaces, in these networks once the band selection is done, the operation is restricted to a single relatively narrow range of channels.

Besides the wide dynamic range, white space links experience uneven fading due to antenna patterns. The fractional bandwidth (FB) for a frequency band, calculated as a ratio of operating bandwidth and the central frequency, determines how wide-band an antenna should be in order to have the same gain over all frequencies with the band. From Table 9.1 we see that white spaces require significantly wider band antennas than GSM and WiFi. Such antennas are hard, if not impossible, to design. Consequently, white space links are highly prone to the effects of imperfect antennas.

To confirm this statement, we deployed a 3 km outdoor non-line-of-sight white space link in a sub-urban area in South Africa shown in Fig. 9.2. The terrain consists of undulating hills and dense foliage. Each of the link nodes consists of a USRP2² software defined radio and a dual core 2.4GHz Pentium PC running GNUradio software. Each radio is equipped with a 8 dBi UHF yagi antenna. One node acts as

² <http://gnuradio.org/>.

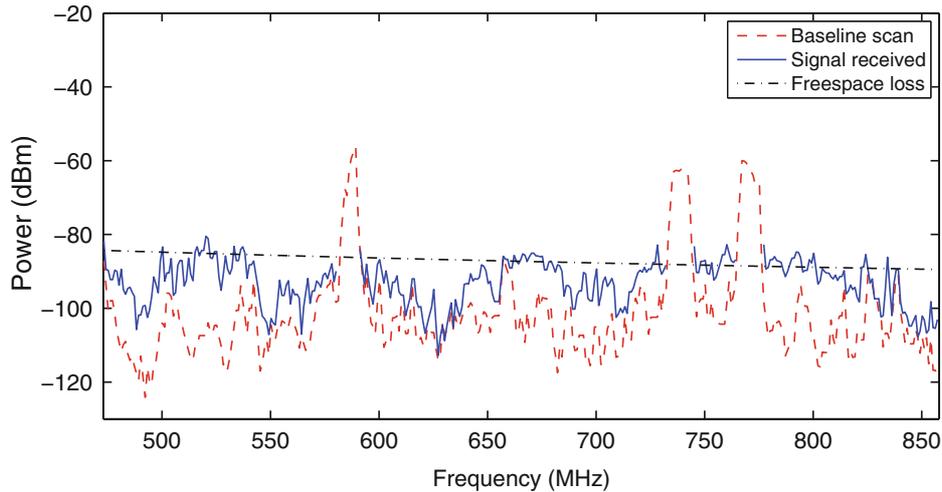


Fig. 9.3 Analysis of received signal strength over the UHF band using 8 dBi yagi antennas at transmitter and receiver. The plots demonstrate that received signal strength is far more dependent on the antenna gain pattern than on attenuation due to free-space loss. This is confirmed by antenna frequency profiles of a number of TV antenna models in Fig. 9.4

a transmitter and sends probes spaced 1 MHz apart over the white space spectrum. The transmitter radio signal is passed through a 1W power amplifier. Another node, the receiver, scans the spectrum in 1 MHz steps both with the transmitter turned off (baseline scan) and with the transmitter sending probes (signal scan).

Figure 9.3 shows the received signal strength across the UHF TV band in South Africa. Three TV stations were detected and probes did not occur at these frequencies. We note that 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 the propagation loss is non-uniform across the UHF band. In order to study the antenna characteristics, we use the WIPL-D antenna modelling package and create a model of the deployed antenna. The results are shown in Fig. 9.4. While an antenna with no surrounding structures has a more predictable gain patterns, when surrounding structures and antenna imperfections, such as bent elements, are introduced, the antenna gain pattern has far less predictability similar to what was seen in our received signal strength measurements. Predicting the type of TV antenna being used or the structures surrounding the antenna is not possible, which necessitates frequency probing in white spaces.

In addition to antenna effects, a part of frequency selectivity may stem from the environment and terrain effects. *Shadowing*, i.e. slow fading due to physical obstacles on the signal path would still be detected and accounted for with frequency probing to detect average channel gain described in Sect. 9.4. Unlike shadowing, detecting the level of *multipath*, which leads to rapid variation of propagation within a channel, requires a more sophisticated channel sounding process but can use the same frequency probes used to detect average channel gain. However, channel allocation

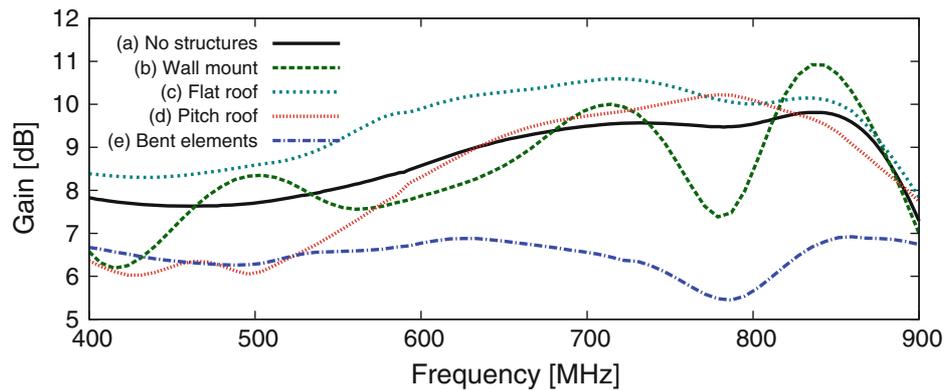


Fig. 9.4 Antenna gain for yagi antenna used in white space link in the outdoor testbed. The plot shows the following scenarios: (a) Antenna with no nearby structures, (b) Antenna mounted on the side of a wall, (c) Antenna mounted on a flat roof, (d) Antenna mounted on a pitch roof, and (e) Antenna with imperfections due to bent elements.

only requires knowledge of the average channel gain of the channel, and in Sect. 9.3 we devise a lightweight channel probing mechanism.

9.2.3 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 [20]. In white spaces, due to the wide dynamic range and antenna effects, the transmission range varies significantly among frequencies in the band (see Table 9.1 and Fig. 9.3). Therefore, *selection of the operating frequency can impact the existence of a link itself*. This further complicates the problem of graph colouring, as now not all colours are equal. Figure 9.5 show one such example where a tradeoff between establishing links and avoiding interference is hard to achieve. In a white space network the color affects the graph structure, thus the existing approaches to frequency assignment are not directly applicable.

9.2.4 Network Architecture

The network scenario that describes the setting in which VillageLink will operate is given in Fig. 9.6. In this paper, we consider wide-area white space networks that consist of individual base stations (BSs), each with a set of associated customer-premises equipment (CPE) clients. We term one such BS with its CPEs *a cell*. A BS

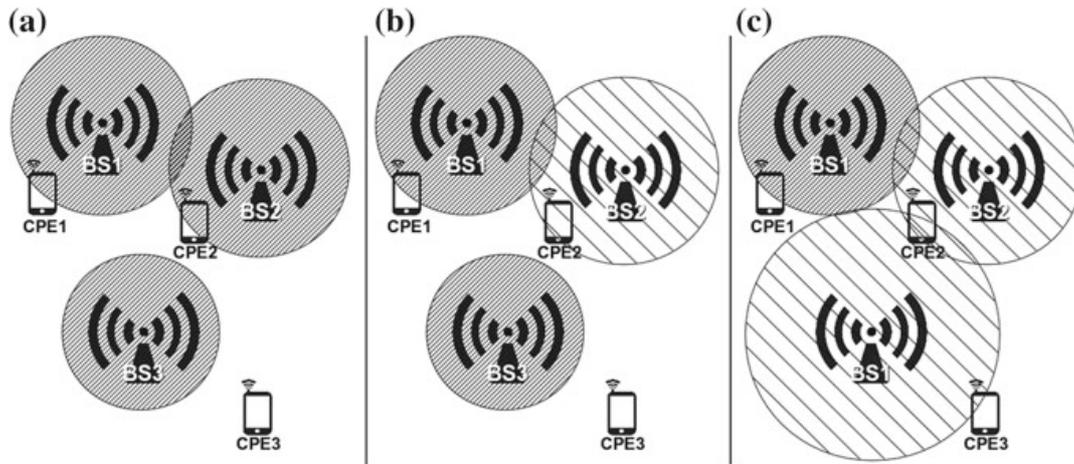


Fig. 9.5 A simple example of the challenges of frequency assignment in white spaces. We want to establish connections from base station (BS_1) to client premises equipment (CPE_1), BS_2 to CPE_2 , and BS_3 to CPE_3 . Two channels represented by different hatching patterns are given. In **a** all links operate on the one of the channels, however there is interference between BS_1 and BS_2 , which may affect clients in the interference area; in addition, due to the channels poor propagation properties BS_3 signal is not reaching CPE_3 . In **b** the interference is resolved by switching the frequency for BS_2 . In **c** a link between CPE_3 and BS_3 is establishment through assignment of a different frequency to BS_3 . However, this frequency propagates further, introducing interference at CPE_2 served by BS_2

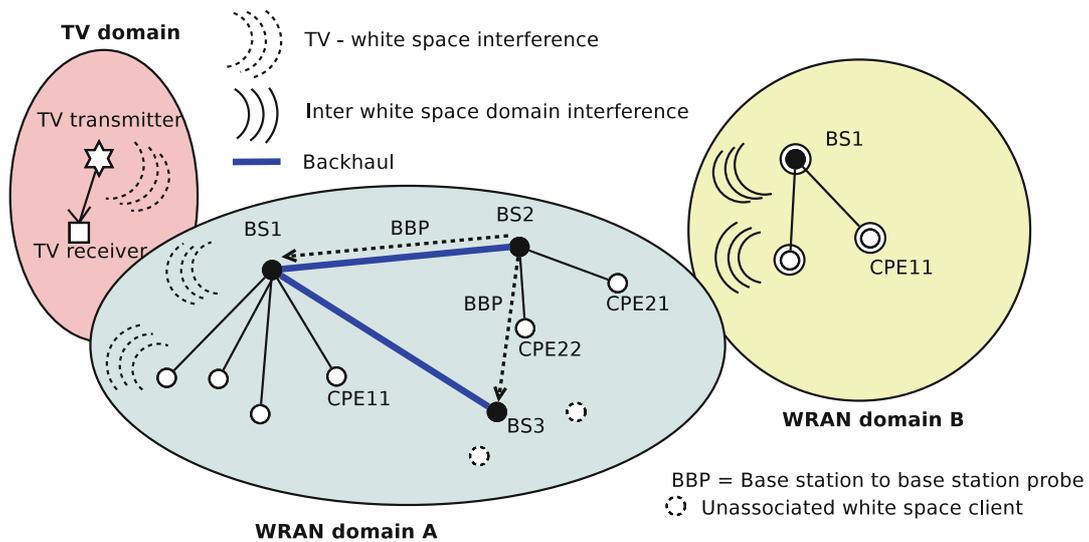


Fig. 9.6 Layout of a targeted white space network showing interference scenarios between television and white spaces, and between white space networks in different domains. White space base stations within the same domain send base station to base station probes (BBPs) to calculate the channel conditions among themselves

and all the CPEs within a cell operate on the same channel³; thus, when considering

³ We envision OFDMA channel sharing among the CPEs of a BS. Such an approach is mandated by IEEE 802.16 and IEEE 802.22 standards. We leave the details of subcarrier allocation as the future work, and in this paper concentrate solely on channel allocation at the BS level.

channel allocation we use “BS” and “cell” interchangeably. All cells that are operated within the same administration are called a WRAN domain. The existence of TV transmission and other 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, as well as supporting structures such as a MAC layer and a frequency probing mechanism, so that the overall network performance within our white space WRAN domain is maximized.

We assume that base stations are connected with a back channel. This can be another white space radio tuned to a common channel that does not interfere with the set of channels available for base station to client connectivity. Moreover, because the amount of control data sent over the back channel is low, a VHF/UHF packet radio, or any other low bandwidth communication technology can be used.

9.3 Medium Access Control Modifications to Support Channel Probing

Channel probing is a necessary tool for propagation evaluation over a wide white space frequency range. Unfortunately, the existing MAC protocols proposed for wide area networks [2, 18, 19] 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. It specifies that the channel may be chosen from the available list by an operator or by a “local routine”. Thus, 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. The details of the technique to calculate the channel conditions using a probe are provided in Sect. 9.4. In this section we outline how we extend the 802.22 MAC protocol to support probing.

The MAC is organised into 160 ms superframes which consist of sixteen 10 ms frames. Each frame is divided into a downstream subframe and an upstream subframe; the size of these subframes depends on the amount of downstream and upstream data that needs to be sent between the base station and the clients. All base stations are equipped with a GPS and are able to synchronize their clocks to within 2ns, which is sufficiently accurate, given the timing requirements of IEEE 802.22. Base stations start their first superframe at the start of a GPS minute cycle; this greatly simplifies inter-base station communication and scheduling. Final synchronization is carried out using the superframe preamble.

The current 802.22 MAC specification already has sophisticated mechanisms to detect other 802.22 WRAN domains using co-existence beacons, move into a time division multiplex mode in cases where base stations are forced to share the same channel, allow clients to scan for and associate with base stations, and carry out ranging between clients and base stations to account for propagation delays. However

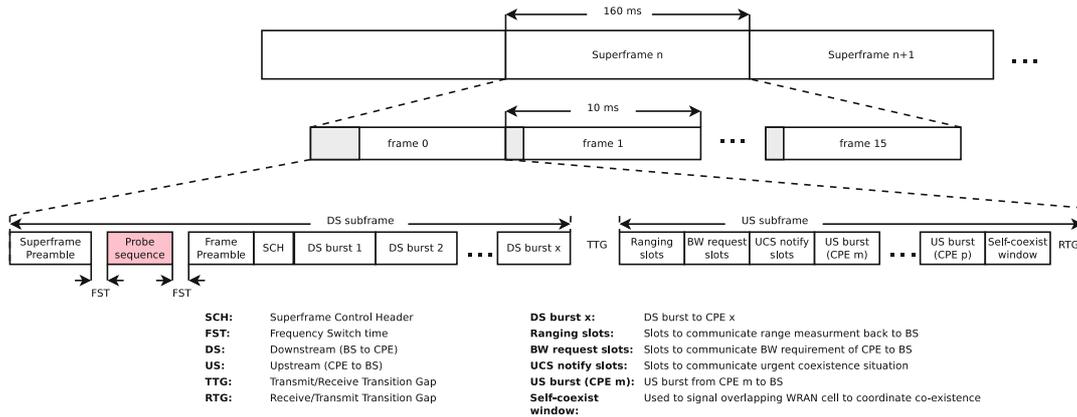


Fig. 9.7 General superframe structure of 802.22 showing additional probe sequence required for frequency profiling. The 1,024 bit probe sequence, in the worst case, increases overhead by 0.44 %; this includes the 10 μ s Frequency Switch Time (FST)

none of these mechanisms allow a base station to discover channel conditions between WRAN cells at all available channels.

Figure 9.7 shows the superframe structure of 802.22 with our modification, an additional probe sequence after the superframe preamble to allow for frequency profiling at different wireless frequencies. The probe sequence uses a 1,024 bit PN sequence modulated using DBPSK with four samples per symbol. The probe sequence incurs low overhead, using an additional 0.44 % of the channel in the worst case. We place the probe sequence after the superframe preamble to ensure we do not break any timing synchronization. Clients make use of the superframe preamble to synchronize any clock offsets. The probe will be transmitted on a probe channel, p_i , where $\mathbf{p} = (p_1, p_2, \dots, p_N)$, a set of N probe channels. The probe channel set is a subset of the complete set of available TV channels, $\mathbf{v} = (v_1, v_2, \dots, v_M)$ with $N < M$ after eliminating non-vacant channels. Each base station in the WRAN domain would first have consulted its spectrum database and scanned all the channels for primary users and other white space domains to ensure availability.

In order to perform frequency profiling between base stations on all available channels, a mechanism is required to coordinate probing timing, channel probe senders and listeners. When a base station is in a probing state, it sends a probe at the beginning of each superframe. It sequentially steps through the full white space TV channel set and only sends a probe if the channel is contained in the probe channel set for that base station. The entire scan takes $160 * M$ ms, where M is the total number of TV channels. If the probe channel is not contained in the probe set, the base station does not send any transmission in the probe sequence slot. Base stations maintain their own probe sets as they may each generate interference to primary users on different sets of frequencies. We chose a mechanism where we step through the entire set of white space channels with interspersed quiet periods on non-vacant channels to avoid needing to maintain full consistency of all probe channels amongst all base stations and associated clients.

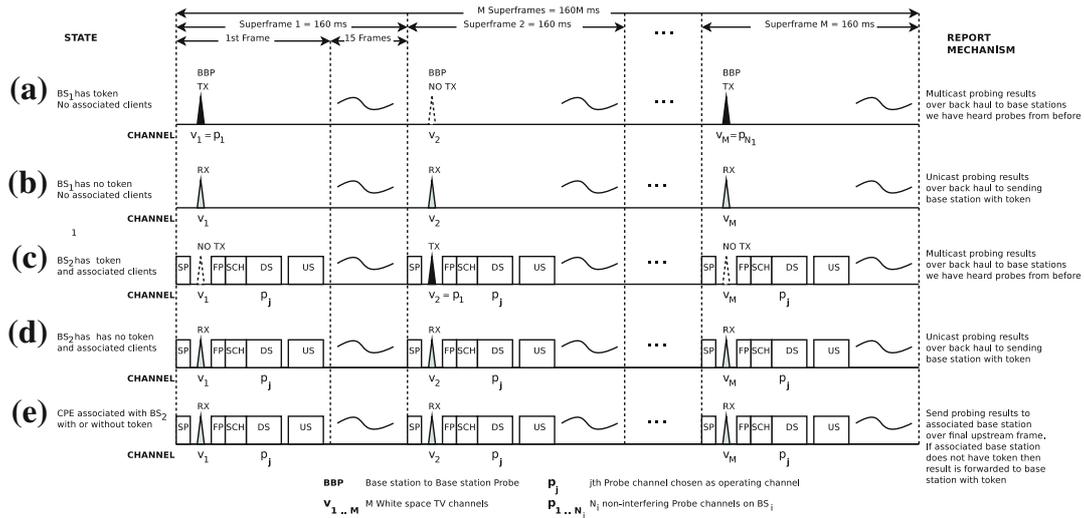


Fig. 9.8 Timing diagram showing the base station probes (BBPs) sent by the base station and received by other base stations and clients. **a** BS₁ has token with no associated clients. **b** BS₁ has no token and no associated clients. **c** BS₂ has token with associated clients. **d** BS₂ has no token and associated clients. **e** CPE associated with BS₂ with or without token

In order to coordinate probe transmission and reception, we propose a token approach in which a base station only transmits probes when it has a probe token and listens for probes when it does not have a token. A breath-first traversal of a spanning tree of the graph is used to ensure that a token traverses the graph of base stations when the back channel forms the edges. Once the base station has finished probing across the TV channel set, it sends the token to the next base station using a traversal algorithm; the probing process is completed when all base stations have used the token. Probes can take place when base stations are in the initialization phase and have not chosen an operating channel, or when they are in an operating phase and are communicating with associated clients and a new base station is added, for example. A request to probe is broadcast by the base station wishing to initiate probing in the domain.

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

Figure 9.8 shows a number of timing diagrams for different base station and client states. Figure 9.8a shows a base station in the initialization phase in which it has a token. In this case only probes are sent by the base station. Figure 9.8c also shows a base station with a probe token but in this case the probes are interspersed within standard 802.22 frames as the base station is actively communicating with clients. Quiet periods, where no probe is sent, are different in Figs. 9.8a and c as they effect primary users on different channels. Figures 9.8b and d show a base station receiving probes and measuring the channel condition between the sending base station and itself. Figure 9.8e shows a client receiving probes from a base station with or without a token. In the case of a client hearing a probe from a base station it is not associated with, the frequency profile matrix is sent back to its associated base station and then forwarded on the back channel to the sending base station with the token.

Once the probing process is completed each BS i has 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 i and each of the BSs j that heard probes from i , and (3) within cells that are served by each of the neighbouring BSs j . We note that power failures are expected in rural areas [21], and should they be experienced, the base station can power up to its previous known state stored in non-volatile memory as the frequency profile matrix, which contains the frequency profile between all base stations at all available channels.

9.4 Channel Profiling

In this section we provide details of the probe that is sent at the beginning of each superframe. This probe allows us to determine the channel gain at each of the available white space channels between base stations and clients. The probe also allows us to determine interference levels between multiple base stations. For determining the optimal channel to use at a base station, we only require a calculation of the overall channel gain at each available frequency. However, once an optimal channel is chosen, the same probe can be used to calculate the detailed channel impulse response for the chosen channel—a process called channel sounding. The channel impulse response can be used for optimal allocation of OFDM sub-carriers to each client connecting to the base station.

9.4.1 Channel Gain Calculation

Each master probe is simply a pseudo-noise (PN) sequence modulated with DBPSK. The client calculates the average power measured, P_{avg} , over the probe listen window.

$$P_{avg}(dB) = 10 \log \left(\frac{1}{N} \sum s(n)^2 \right) + CF \quad (9.2)$$

where $s(n)$ is the signal received, n is the sample number, N is the total number of samples and CF is the correction factor, which is calculated by calibrating the receiver. Average power is a low complexity ($O(N)$) calculation and the cognitive radio can carry this out in real time using a cumulative average.

The SNR of the probe can then be calculated by using a previously known noise level N_{avg} from an initial scan when no probes are present.

$$SNR(dB) = P_{avg}(dB) - N_{avg}(dB) \quad (9.3)$$

Figure 9.3 shows the result of this channel gain calculation across the entire UHF TV band. The average noise shown in the Figure is measured when no probes are sent and the average power is measured in 1 MHz steps when probes are sent from the base station.

From the measured SNR we can extract the channel gain:

$$H = \frac{SNR \cdot N_0 \cdot W}{P} \quad (9.4)$$

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

The set of channel gains at each probe frequency are used for the channel allocation algorithm described in Sect. 9.5.

9.4.2 Channel Sounding

Once a channel is selected using the channel allocation algorithm, the system can carry out more detailed analysis of the channel. This is done by analyzing the channel impulse response at a specific frequency using the stored received probe sequence to calculate the channel gain at each available white space frequency.

The discrete channel impulse response is calculated using a correlation-based channel sounder. The received signal, $r[n]$, is given by the convolution of the channel, $h[n]$, and the transmitted signal, $s[n]$, where n is the sample number:

$$r[n] = s[n] * h[n] \quad (9.5)$$

In our case, we need to estimate $h[n]$ from the received signal $r[n]$ and a known signal $s[n]$ using deconvolution:

$$h[n] \approx \frac{1}{P_s} r[n] * s[-n]. \quad (9.6)$$

where P_s is the power spectral density of the input signal $s[n]$. As we are using digital signals, it is reasonable to assume that the input signal is approximately flat or

$P_s = |S(n)|^2$ is a known constant. Once $h(n)$ is known we can calculate the Fourier transform of $h(n)$ to find the frequency domain response, $H(n)$.

In order to achieve maximum discrimination in the correlator, a DBPSK modulated PN sequence is chosen such that the Index of Discrimination (ID) is high. ID is defined as the ratio of the first autocorrelation peak to the nearest neighboring autocorrelation peak. The receiver has prior knowledge of the PN sequence that was sent by the transmitter and can carry out the deconvolution with the received signal in order to calculate the channel impulse response. The frequency response of the channel is determined by Fourier transforming the channel impulse response.

The minimum multipath delay resolution that is possible is given by $\delta\tau = \frac{2}{\alpha}$, where α is rate of the PN sequence also known as the chipping rate. The minimum path difference, D_{pd} , that can be resolved between multipath components is given by $D_{pd} = c\alpha$. It is important to determine the level of multipath in a channel as increased levels of multipath cause less coherence in the frequency response of the channel. Although channel equalization can be achieved by adjusting the magnitude of the OFDM sub-carriers at the transmitter, selecting sub-carriers for each client is a complex problem due to different channel conditions between a base station and each of it's associated clients.

We store ten PN sequence lengths worth of received data in memory and carry out the deconvolution over this data set. This allows us to store ten different channel impulse responses for each repeated PN sequence and carry out transient analysis or averaging.

9.4.2.1 Practical Considerations

We use a chipping rate of 1 and 2 MHz for our experiments. We found that higher chipping rates do not perform well on the USRP radios. This translates to a minimum multipath delay resolution of 2 and 1 μ s or a minimum path difference of 600 and 300 m respectively. From previous studies [4] rural open areas typically have delay spread of up to 1.3 μ s or up to 11.6 μ s if there are deep canyons or mountains. Suburban areas usually do not have delay spread over 2.4 μ s and urban areas have delay spread up to 4.6 μ s with many reflections. From this study our chipping rate should be sufficient to discriminate multipath in most environments.

To further improve the signal to noise ratio of our channel sounder, we sample at 4 MHz at the receiver. This receiver rate, between 2 and 4 times the PN sequence chip rate, creates an oversampled signal that improves the SNR. We also improve the SNR by using a long PN sequence (1,024 bits) taking multiple snapshots of the received PN sequence and calculating the average.

9.4.2.2 Results

Figure 9.9 shows the results for channel sounding at 530 MHz and a lower chip rate of 1 MHz. This is shown for a single deconvolution snapshot over one PN sequence.

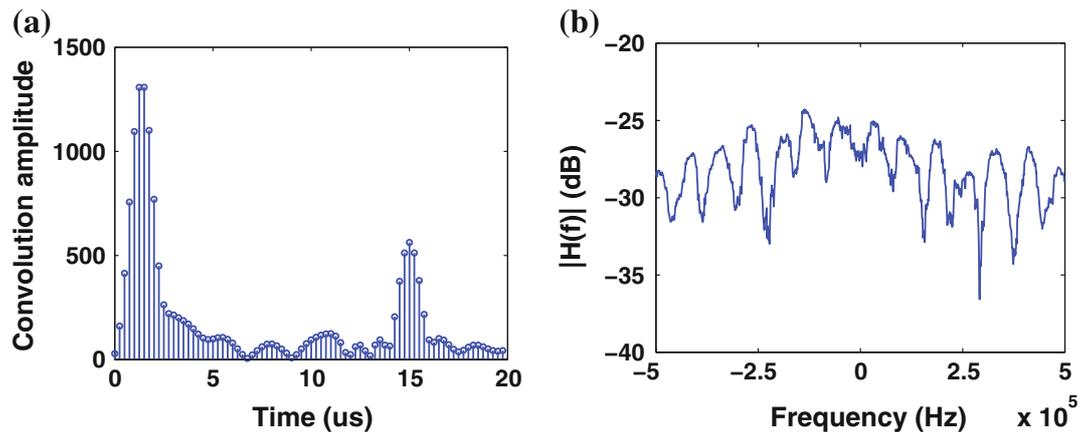


Fig. 9.9 Channel sounding at 530 MHz using 1 MHz chipping rate and oversampling rate of 4. **a** Is the channel impulse response of the channel and **b** is the frequency response of the channel. This plot shows the result for deconvolution of a single received PN sequence to determine the channel impulse response

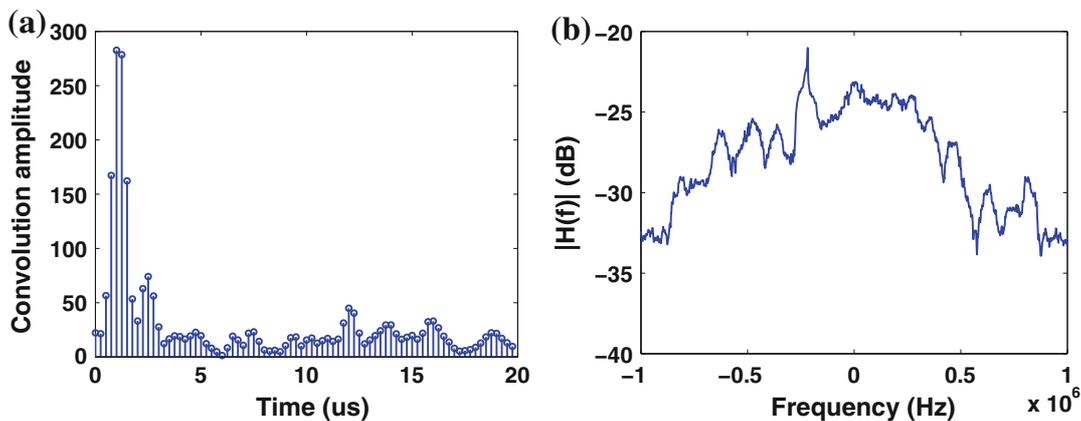


Fig. 9.10 Channel sounding at 530 MHz using 2 MHz chipping rate and oversampling rate of 2. **a** Is the channel impulse response of the channel and **b** is the frequency response of the channel. This plot shows the result for averaging the deconvolution of 10 received PN sequences to determine the channel impulse response

There is a very clear strong echo at 15 μs ; there is also an echo at approximately 2 μs but it is hard to distinguish from the main direct path as the chipping rate is too low. Ideally you want the chipping rate resolution to be at least half that of the time to echo. The frequency response shown in Fig. 9.9b shows some deep fades due to the large echo present in the channel.

In Fig. 9.10, the chipping rate is increased to 2 MHz and we generate an average of 10 channel impulse responses. The first echo at 2 μs is now very visible, however the 15 μs echo is no longer visible. It is possible that this echo was temporary due to a change in the environment, such as a moving truck. Fewer deep fades are visible in the frequency response.

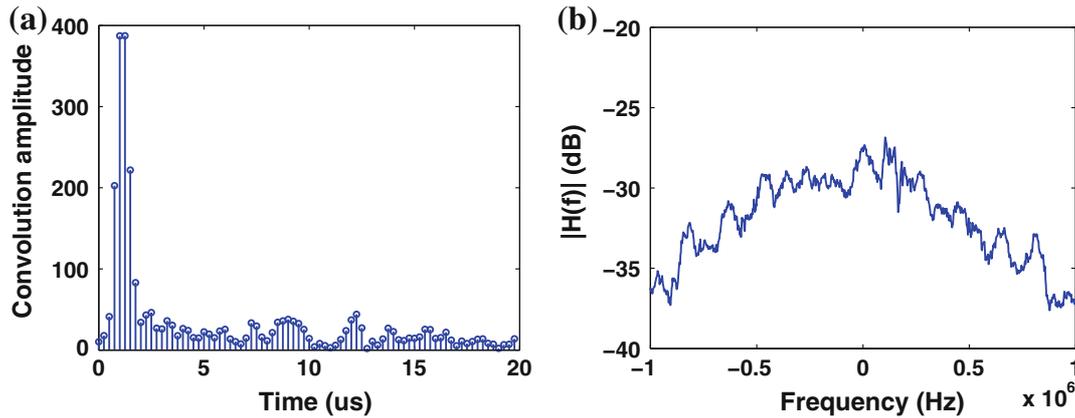


Fig. 9.11 Channel sounding at 834 MHz using 2 MHz chipping rate and oversampling rate of 2. **a** Is the channel impulse response of the channel and **b** is the frequency response of the channel. This plot shows the result for averaging the deconvolution of 10 received PN sequences to determine the channel impulse response

Figure 9.11 shows the channel impulse response for 834 MHz at the 2 MHz chipping rate. There are fewer echoes than in the lower UHF frequency plots, and as a result this frequency response is the most coherent. Environmental terrain has different absorption levels at different frequencies and this specific sub-urban environment had higher absorption levels for multipath signals at higher frequencies.

In the rest of the chapter, we only make use of the channel gain calculations for optimal channel selection and we leave optimal allocation of sub-carriers using the channel impulse response for future work.

9.5 Channel Allocation

In this section we devise a distributed channel allocation algorithm that uses information obtained through frequency profiling (Sect. 9.3) and does not incur channel switching overhead typical for other allocation schemes. Our approach is based on the annealed Gibbs sampler, a technique that can help us minimize a target function in a distributed way. In the next subsection we present the basics of Gibbs sampling. An interested reader can find more details about Gibbs sampling in [1]. We then cast our problem to the Gibbsian framework and sketch the channel allocation algorithm.

9.5.1 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 $j = 1, \dots, k$, 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 $x_1^j, \dots, x_N^j; j \in [1, \dots, k]$ samples from the joint distribution p .

We can solve the channel allocation problem through Gibbs sampling, if we obtain the samples from a multivariate probability distribution that:

1. Is related to overall network performance.
2. Depends on the selected operating channel of each of the base stations.
3. Isolates the impact of each of the base stations on the total optimization function.
4. Can be calculated in a distributed way and sampled independently at each of the base stations.
5. Favors states that lead to maximum network performance.

In the following section we develop a network performance metric that can be used as a basis for a probability distribution that satisfies the above demands.

9.5.2 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:

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

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:

$$C(\mathbf{c}) = \sum_i C_i(c_i) = \sum_i W \log(1 + SINR_i(c_i)) \quad (9.7)$$

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 using Gibbs sampling. Centralized optimization using known polynomial complexity techniques, such as linear programming, is not directly applicable either, since the target sum involves non-linear factors and discrete variables.

One of the ways to circumvent this is to revert to a tighter problem formulation that prevents interfering base stations from concurrent transmission [11]. While this can be enforced in a network that employs carrier sensing and collision avoidance, in our setting long distances between base stations render such coordination inefficient [18]. In addition, allowing some interference often yields more capacity than restricting concurrent transmissions [15]. Another approach is to modify the optimization function and instead of maximizing capacity concentrate on minimizing total network interference [9, 17]. This approach is attractive for networks, such as WiFi, where these two goals are essentially interchangeable. In a white space setting, 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—*Cumulative interference plus noise to signal ratio (CINSR)*—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. Compared to metrics such as the total capacity or the overall level of interference, CINSR takes into account the frequency diversity that exists in white space networks, and allows distributed performance optimization with Gibbs sampling:

$$CINSR(\mathbf{c}) = \sum_{i=1}^N \frac{1}{SINR_i(c_i)} \quad (9.8)$$

$$= \sum_{i=1}^N \frac{N_0 W + \sum_{j=1..N, j \neq i} ch(i, j) PH_{ji}(c_j)}{PH_i(c_i)} \quad (9.9)$$

The first term in the numerator within the above sum is the thermal noise (a product of the channel width W and the noise constant N_0), whereas the second term is the sum of interference experienced at cell i , and originating from all other base stations

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 base station 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 that transmits with power P . The average channel gain towards the clients 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*:

$$CINSR_i(\mathbf{c}) = \frac{N_0W}{PH_i(c_i)} \quad (9.10)$$

$$+ \sum_{j \neq i} ch(i, j) \left(\frac{PH_{ij}(c_i)}{PH_i(c_i)} + \frac{PH_{ji}(c_i)}{PH_j(c_i)} \right) \quad (9.11)$$

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 Sect. 9.3.

9.5.3 The Gibbs Distribution

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

$$\pi(\mathbf{c}) = \frac{e^{-CINSR(\mathbf{c})/T}}{\sum_{\mathbf{c}' \in c^N} e^{-CINSR(\mathbf{c}')/T}} \quad (9.12)$$

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 Sect. 9.5.1.

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 $\pi_i(\mathbf{c})$:

$$\pi_i(\mathbf{c}) = \frac{e^{-CINSR_i(c_i, (c_j)_{j \neq i})/T}}{\sum_{\mathbf{c}' \in c^N} e^{-CINSR_i(c'_i, (c'_j)_{j \neq i})/T}} \quad (9.13)$$

and transitions to the sampled local state, converging to the stationary distribution $\pi(\mathbf{c})$ (see Sect. 9.5.5).

Distribution $\pi(\mathbf{c})$ highly favors low $CINSR$ states when the temperature T 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.

9.5.4 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 [9, 14], 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 base stations 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.

9.5.5 Algorithm Convergence

Convergence of a Gibbs sampler, and its annealed version, is a well researched topic [1]. 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 (Eq. 9.12) 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:

$$\mathcal{E}(\mathbf{s}) = \sum_{\Gamma} V_\Gamma(\mathbf{s}) \quad (9.14)$$

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 (p. 260 in [1]), and we proceed with showing that the function that we use to construct the Gibbs distribution in Eq. 9.12— $CINSR(c)$ 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 Eq. 9.14 and $CINSR$ can be used as the energy function for Gibbs sampling:

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

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 variation⁴ 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 (Eq. 9.12) is the invariant probability measure (Example 6.5, p. 288 in [1]). Note that direct sampling of the capacity (Eq. 9.7) 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, p. 311 in [1]). Conditions that the cooling schedule has to satisfy in order for convergence to happen can be found in [5].

9.6 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 [16]. 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 Sect. 9.2.2, was performed on a 3 km outdoor link.

9.6.1 Simulation Setup

For a comprehensive evaluation of the channel allocation algorithm, we rely on a Matlab-based custom simulator. The simulator allows us to scale our experiments over a number of cells, and to model different network layouts. We explicitly take into account high variability of signal propagation in the white space band by modeling propagation with the Friis transmission equation:

$$P_r = P_t + G_t + G_r + 20 \log \left(\frac{\lambda}{4\pi R} \right)$$

where λ , R , P_r , P_t , G_t and G_r are the wavelength, distance between antennas, received power, transmitted power, transmitter antenna gain, and receiver antenna

⁴ Convergence in variation describes convergence of an array of samples to a probability distribution and is defined in [1], p. 128.

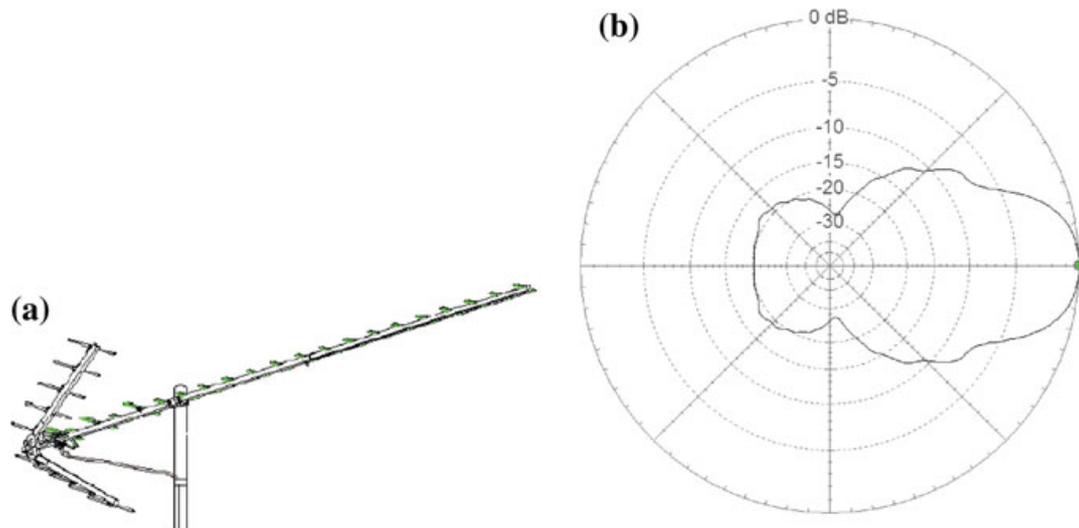


Fig. 9.12 Wineguard PR9032 UHF Yagi/corner reflector antenna used as a base station antenna in our evaluation. Showing **a** the antenna design and **b** its radiation pattern seen from the top of the antenna

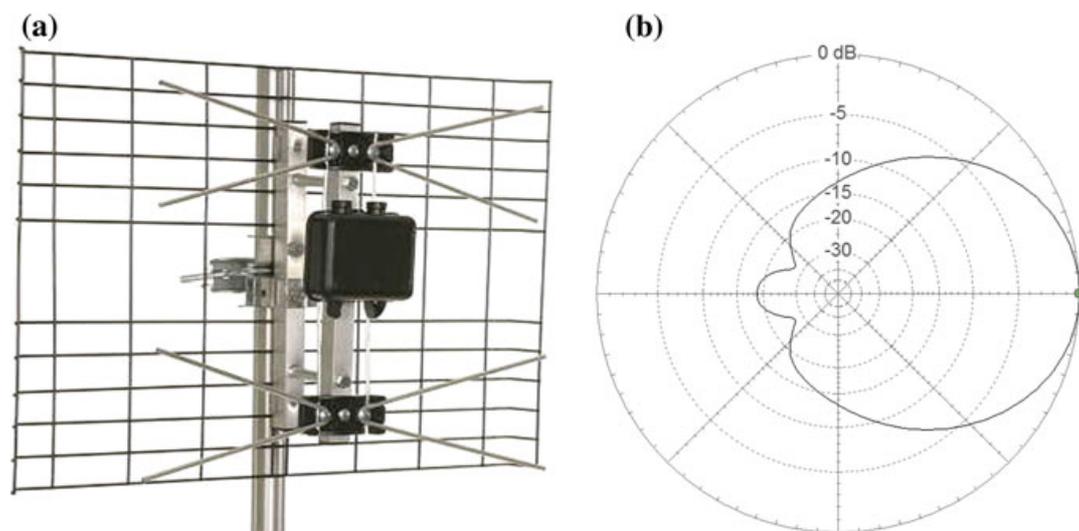


Fig. 9.13 AntennasDirect DB-2 2-Bay UHF antenna; one of the client antennas used for the evaluation. Showing the antenna in **(a)** and its radiation pattern seen from the top of the antenna in **(b)**

gain, respectively. Antenna gains depend on specific devices used and their orientations. 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 (Fig. 9.3, Sect. 9.2), thus we model antenna effects in detail.

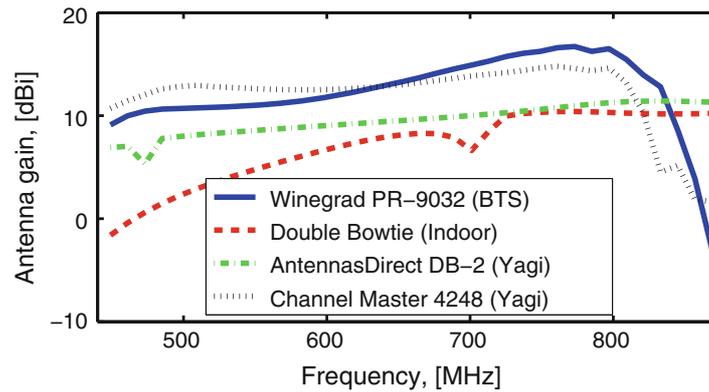


Fig. 9.14 Antenna profiles of four of the antennas used in our evaluation. One of the profiles, Wineguard PR-9032, corresponds to the BS antenna and, the other three, to client antennas

We use publicly available antenna models⁵ and the Numerical Electromagnetics Code (NEC)⁶ antenna modeling software to examine propagation over different frequencies with different antennas. Figures 9.12b and 9.13b show the radiation patterns seen from the center frequency (598 MHz) of the white space band for two different antennas. In Fig. 9.14 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 Fig. 9.14. Thus, in the simulations we use the antenna pattern shape of the center frequency to account for antenna orientation, and we use the full gain over frequency diversity.

All base stations in our simulations use the Yagi antenna from Fig. 9.12, 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 15 dBi to simple indoor loop antennas with a gain of 3 dBi.

We run our experiments over a white space band from 443 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 100 km \times 100 km 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 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.

⁵ <http://www.hdtvprimer.com/ANTENNAS/comparing.html>.

⁶ www.nec2.org.

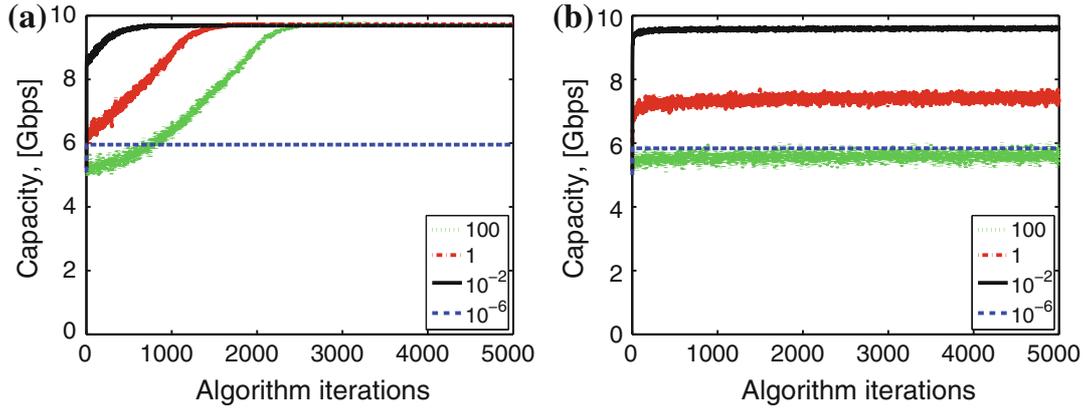


Fig. 9.15 Algorithm convergence with the (a) exponential, and (b) logarithmic cooling schedule. Different lines correspond to different starting temperatures

9.6.2 Channel Allocation Convergence

We simulate Algorithm 1 behavior in a network of ten base stations and five white space channels that are available for communication. We are interested in the algorithm convergence under different Gibbs sampling parameters. We experiment with two common cooling schedules:

- Logarithmic: $T = T_0 / \log(t + 2)$, proposed in [3].
- Exponential: $T = T_0 \alpha^t$, proposed in [10].

where T denotes the temperature at time t , T_0 is the starting temperature, and α is a real number between zero and one; we empirically find value 0.995 to work well in our experiments.

The selection of the starting temperature is important for proper annealing. In Fig. 9.15 we plot total network capacity achieved with the two schedules and four different starting temperatures for each. Each point in the graphs is an average over 100 runs. 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 the starting temperature, between the convergence time and the assurance that the optimal value will be found. In the rest of the evaluation section we fix the starting temperature to 1, a value that allows full exploration of the solution space and converges in a reasonable amount of time.

We observe much faster convergence with the exponential schedule, that converged in all but one case ($T_0 = 10^{-6}$). The logarithmic schedule did not converge in 5,000 iterations for $T_0 = 100$ and $T_0 = 1$. In the rest of this section we rely exclusively on the exponential schedule.

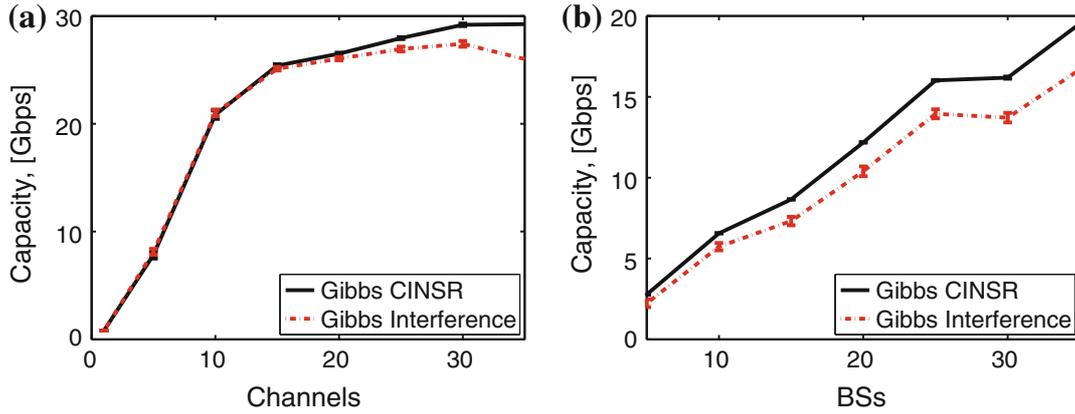


Fig. 9.16 Comparison of the total network capacity achieved with CINSR and Interference metrics. We simulate under-provisioned and over-provisioned number of channels with respect to the number of base stations in the network. **a** Channel under-provisioning. **b** Channel over-provisioning

9.6.3 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—which is often used as a metric in channel allocation algorithms [9, 17].

The total network interference and noise is defined as:

$$I(\mathbf{c}) = \sum_{i=1}^N \left(N_0 W + \sum_{j=1..N, j \neq i} ch(i, j) PH_{ij}(c_i) \right) \quad (9.15)$$

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

$$I_i(\mathbf{c}) = N_0 W + \sum_{j=1..N, j \neq i} ch(i, j) (PH_{ij}(c_i) + PH_{ji}(c_i)) \quad (9.16)$$

We modify the Gibbs distribution (Eq. 9.12) to include $I(\mathbf{c})$ instead of $CINSR(\mathbf{c})$, and the local Gibbs distribution (Eq. 9.13) to include $I_i(\mathbf{c})$ instead of $CINSR_i(\mathbf{c})$. 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 not account for the balance between well propagating channels that are preferred by the CPEs and inferior channels that minimize inter-cell interference.

9.6.3.1 Channel Under-Provisioning

In the first scenario we simulate a network with a number of contending BSs higher than the number of available channels. This can be the case in the urban developed world, for example. We put 50 cells in the same $100 \text{ km} \times 100 \text{ km}$ region. We experiment with a varying number of available channels. The total network capacity is plotted in Fig. 9.16a. Each point represents an average value of 20 runs of the algorithm with a different metric, *Gibbs CINSR* or *Gibbs Interference*, over the same topology.

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 Fig. 9.16a 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 Eq. 9.15, this version of the Gibbs sampler results in a channel allocation that delivers less capacity than the version that uses *CINSR*.

9.6.3.2 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 Fig. 9.16b. Each point represents an average value of 20 runs of the algorithm (*Gibbs CINSR* or *Gibbs Interference*) over the same topology.

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 $H_i(c_i)$ *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 areas where the channel availability is high* [6].

9.6.4 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:

- *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 [14].

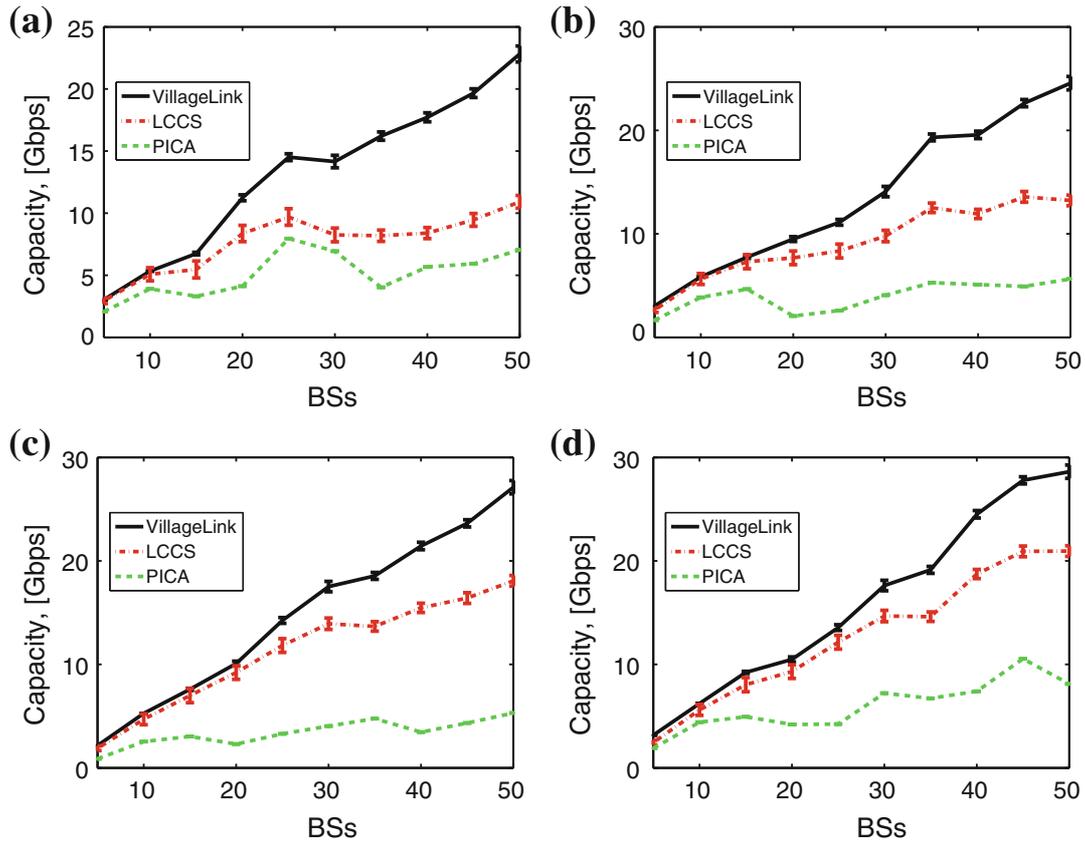


Fig. 9.17 Total network capacity with varying number of channels and base stations. **a** 10 available channels. **b** 15 available channels. **c** 20 available channels. **d** 25 available channels

- *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 ($\arg \max_{c_i} H_i(c_i)$).

These heuristics optimize a non-submodular capacity function in a greedy manner, therefore may settle for a solution that is arbitrarily far from the optimal. VillageLink’s convergence to the states of minimum *CINSR* is proven in Sect. 9.5.5. 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.

9.6.4.1 Total Network Capacity

In Fig. 9.17 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, 15, 20 and 25 available channels. We plot average values and two standard deviations (represented by error bars) for each data point.

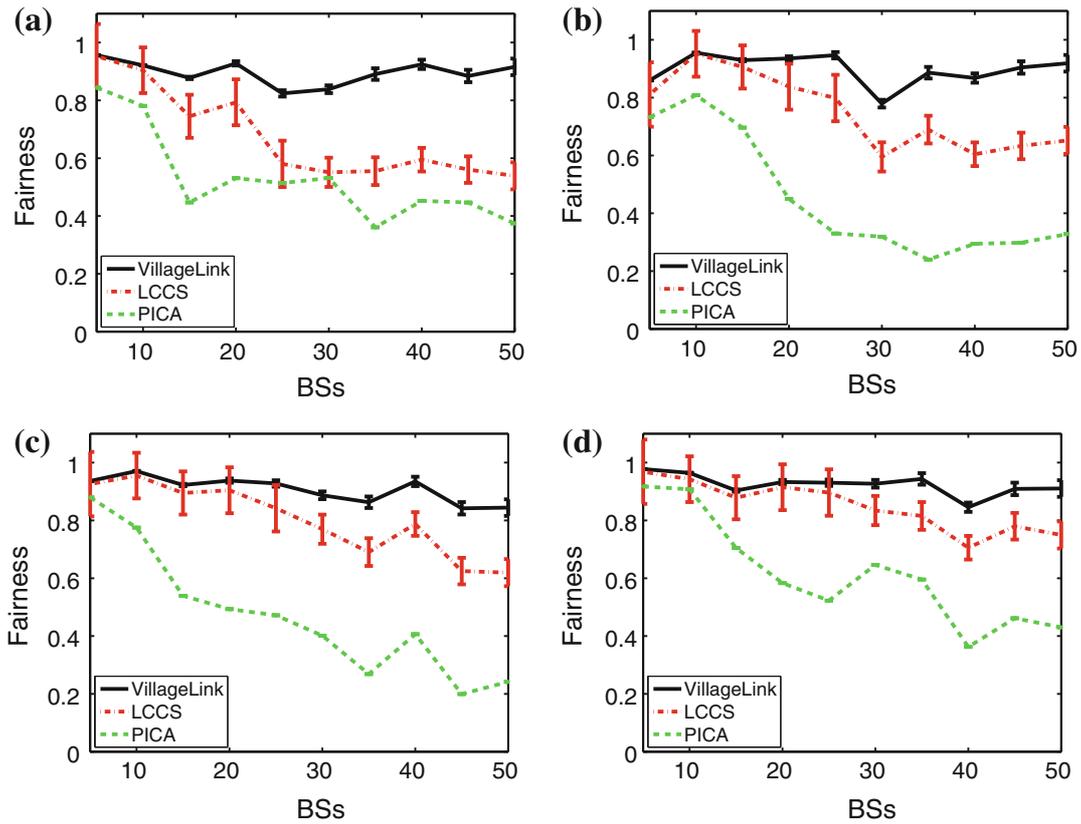


Fig. 9.18 Fairness with varying number of channels and base stations (the closer the fairness index value is to one—the better). **a** 10 available channels. **b** 15 available channels. **c** 20 available channels. **d** 25 available channels

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 and 50 BSs—15 channels, VillageLink delivers twice as much capacity as the next best alternative, LCCS. A comprehensive comparison of LCCS and PICA could unravel the importance of two conflicting goals in channel allocation: minimizing interference and maximizing intra-cell capacity in isolation, and is left for future work.

9.6.4.2 Fairness

In Fig. 9.18 we plot the Jain fairness index [8] 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.

9.7 Related Work

Efforts to provide broadband connectivity to remote rural regions with low-cost unlicensed options, such as modified WiFi, have been proposed in the last decade [18, 19]. While numerous rural WiFi deployments provide useful general guidelines for wide-area coverage, the propagation characteristics in white spaces are drastically different than in the WiFi bands, and networking protocols have to be reconsidered for the new spectrum. IEEE 802.22 [2] is a standardized protocol for wide area white space coverage. In VillageLink we embrace the 802.22 frame structure, and augment it with novel channel probing and operating frequency selection mechanisms.

Previous work related to channel assignment in wireless networks usually casts the problem of channel assignment as an NP-hard graph coloring problem [14]. Numerous heuristics have been proposed to provide an approximate solution ([20] and references therein). Ma and Tsang [11] recognize the channel heterogeneity in the case of wide bands and propose an integer linear programming solution for the frequency allocation problem. However, they restrict frequency reuse to well defined interference domains, thus no two BSs are allowed to transmit at the same time if they interfere. Motivated by [15], we rely on a more sophisticated representation of interference—we directly measure its impact through probing and account for any interference during the allocation process.

The Gibbs sampling, under this name, was first proposed in 1984 for image manipulation [3]. Its applicability to distributed channel allocation, client association and power control in wireless networks has been realized over the last twenty years [9, 13]. VillageLink differs from these by two important properties. First, we propose a novel network performance metric called *CINSR*, that takes into account frequency dependence of both useful signal transmissions and interference. Second, our algorithm does not require subsequent channel switching and environment sensing after each local decision is made. Rather, only control information has to be exchanged among neighboring nodes, and once the algorithm terminates only a single channel switch is made per node.

9.8 Conclusion

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. In this work 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.

White space networks are largely unexplored, and their straightforward implementation might prove difficult due to unique characteristic they exhibit. For example,

experiments we performed on VillageLink demonstrate the necessity of careful channel allocation in white spaces even when the number of devices operating in the same interference domain is low, which is a stark contrast to WiFi networks. Our work examines only one aspect of network adaptation. The complex nature of signal propagation over a wide frequency band opens up new possibilities for protocol design and further refinement of channel access in white spaces.

References

1. Bremaud, P.: Markov Chains, Gibbs Fields. Monte Carlo Simulation and Queues. Springer, New York (1999)
2. Cordeiro, C., Challapali, K., Birru, D., Manor, B., Diego, S.: IEEE 802.22: an introduction to the first wireless standard based on cognitive radios. *J. commun.* **1**(1), 38–47 (2006)
3. Geman, S., Geman, D.: Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721–741 (1984)
4. Greenstein, L.J., Erceg, V., Yeh, Y.S., Clark, M.V.: A new path-gain/delay-spread propagation model for digital cellular channels. *IEEE Trans. Veh. Technol.* **46**(2), 477–485 (1997)
5. Hajek, B.: Cooling schedules for optimal annealing. *Math. Oper. Res.* **13**, 311–319 (1988)
6. Harrison, K., Mishra, S. M., Saha, A.: How much white-space capacity is there? In: *DySpan'10*, Singapore (2010)
7. ITU World Telecommunication/ICT Indicators Database. International Telecommunication Union (2013)
8. Jain, R., Chiu, D.M., Hawe, W.R.: A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System. Digital Equipment Corporation, Eastern Research Laboratory (1984)
9. Kauffmann, B., Baccelli, F., Chaintreau, A., Mhatre, V., Papagiannaki, K., Diot, C.: Measurement-based self organization of interfering 802.11 wireless access networks. In: *INFOCOM'07*, Anchorage (2007)
10. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by simulated annealing. *Science* **220**, 671–680 (1983)
11. Ma, M., Tsang, D.H.K.: Joint design of spectrum sharing and routing with channel heterogeneity in cognitive radio networks. *Phys. Commun.* **2**(1–2), 127–137 (2009)
12. Masonta, M., Johnson, D.L., Mzyece, M.: The white space opportunity in Southern Africa: measurements with meraka cognitive radio platform. In: *e-Infrastructure and e-Services for Developing Countries*, Volume 92 of Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, pp. 64–73. Springer, Heidelberg (2012)
13. Mhatre, V., Papagiannaki, K., Baccelli, F.: Interference mitigation through power control in high density 802.11 WLANs. In: *INFOCOM'07*, Anchorage (2007)
14. Mishra, S., Hwang, J., Filippini, D., Moazzami, R., Subramanian, L., Du, T.: Economic analysis of networking technologies for rural developing regions. *Lect. Notes Comput. Sci.* **3828**, 184–194 (2005)
15. Moscibroda, T., Wattenhofer, R., Weber, Y.: Protocol design beyond graph-based models. In: *HotNets'06*, Irvine (2006)
16. Nychis, G., Hottelier, T., Yang, Z., Seshan, S., Steenkiste, P.: Enabling MAC protocol implementations on software-defined radios. In: *NSDI'09*, Boston (2009)
17. Ramachandran, K.N., Belding, E.M., Almeroth, K.C., Buddhikot, M.M.: Interference-aware channel assignment in multi-radio wireless mesh networks. In: *INFOCOM'06*, Barcelona (2006)

18. Raman, B., Chebrolu, K.: Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks. In: *MobiCom'05, Cologne* (2005)
19. Sheth, A., Nedeveschi, S., Patra, R., Surana, S., Subramanian, L., Brewer, E.: *WiLDNet: design and implementation of high performance WiFi based long distance networks*. In: *NSDI, Anchorage* (2007)
20. Stojmenovic, I.: *Handbook of Wireless Networks and Mobile Computing*. Wiley-Interscience, New York (2002)
21. Surana, S., Patra, R., Nedeveschi, S., Ramos, M., Subramanian, L., Ben-David, Y., Brewer, E.: *Beyond pilots: Keeping rural wireless networks alive*. In: *NSDI, San Francisco* (2008)
22. Sydell, L.: *FCC Eyes Broadband For Indian Reservations*. <http://www.npr.org/templates/story/story.php?storyId=128004928> (2010)