Mobile Access of Wide-Spectrum Networks: Design, Deployment and Experimental Evaluation

Anastasios Giannoulis* Paul Patras† Edward W. Knightly*
*Dept. of Electrical and Computer Engineering, Rice University, Houston, TX
†Hamilton Institute, National University of Ireland, Maynooth, Ireland

***Abstract*—Wireless networks increasingly utilize diverse spectral bands that exhibit vast differences in both transmission range and usage. In this work, we present MAWS (Mobile Access of Wide-Spectrum Networks), the first scheme designed for mobile clients to evaluate and select both APs and spectral bands in wide-spectrum networks. Because of the potentially vast number of spectrum and AP options, scanning may be prohibitive. Consequently, our key technique is for clients to infer channel quality and spectral usage for their current location and bands using limited measurements collected in other bands and at other locations. We experimentally evaluate MAWS via a wide-spectrum network that we deploy, a testbed providing access to four bands at 700 MHz, 900 MHz, 2.4 GHz and 5 GHz. To the best of our knowledge, the spectrum of these bands is the widest to be spanned to date by a single operational access network. A key finding of our evaluation is that under a diverse set of operating conditions, mobile clients can accurately predict their performance without a direct measurement at their current location and spectral bands.**

I. Introduction

Wireless networks operating in unlicensed spectrum can now utilize frequency bands ranging from 512 MHz (DTV white spaces) to 5.845 GHz, bands that exhibit vast differences in both transmission range and available airtime. Joint use of multiple diverse bands will therefore provide future network operators with flexibility in coverage provisioning, capacity planning and interference management.

Mobile access of such wide-spectrum networks introduces two key challenges. First, mobile clients must assess and select both APs and spectrum in a timely and efficient manner. Moreover, in wide-spectrum networks, the number of association options, i.e., AP-channel pairs, is significantly higher than in single-band networks. Second, these association options may result in significant differences in client performance. For example, links operating in lower frequencies may offer higher channel quality and lower handoff rate (due to reduced attenuation and increased coverage), yet they are also subject to increased interference due to greater transmission range. Thus, mobile clients must account for multiple conflicting factors in selecting the association option that best meets their individual performance objectives.

Prior work can be classified into three categories: *(i)* Fixed band prioritization is a simple mechanism for mobile clients to select between independent networks operating in dissimilar bands (e.g., smart-phone preference of WLAN to 3G; see also [1]). However, it does not account for the spatial, spectral and temporal variations in channel quality and usage. *(ii)* Scanning dynamically prioritizes association options, e.g., in single-band scenarios spanning WLAN [2] to vehicular WiFi [3]. Unfortunately, scanning either sacrifices airtime that increases with the number of bands or, when additional radios are employed for scanning, increases the power consumption of the energy-constrained mobile clients. *(iii)* Analysis of historical data can also be used to dynamically prioritize association options in mobile access of single-band networks [4] and non-mobile access of multi-band networks [5]. Unfortunately, [4], [5] require measurements at all frequencies and locations, which is prohibitive in wide-spectrum networks.

In this work, we present the following two contributions. First, we propose MAWS (Mobile Access of Wide-Spectrum Networks), the first scheme designed for evaluation and selection of association options by mobile clients of wide-spectrum networks. MAWS is a client-side access solution that considers individualized throughput and delay performance objectives. In contrast to prior work, MAWS clients dynamically prioritize their association options without exclusively employing scanning or historical data for selection of APs and spectrum.

The key technique in MAWS is a method for mobile clients to infer channel quality and spectral usage at their current location without taking a measurement there. To achieve this, MAWS clients first employ limited-rate scanning to measure the spectral usage, and channel quality in as few as two different bands at limited locations. Then, MAWS infers channel quality and spectral usage for the remaining spectrum and space. To estimate *channel quality* in alternate bands and locations than those measured, we introduce cross-spectral and -spatial inference methods that couple the limited measurements with coarse-grained propagation models. To estimate the *usage* of a frequency, MAWS exploits any spatial and temporal correlation of a frequency's usage and calculates a weighted average of its measured usage at locations of prior visits. Each weight is related to the distance to these locations as well as to an estimate of the interference range in the respective band. Finally, the inferred metrics are coupled with estimates of the handoff rate under each association option to drive the calculation of throughput and delay predictions.

This research was supported by the NSF (CNS-1012831 and CNS-1126478 grants) and the European Commission (FP7 grant agreement n. 257263 – FLAVIA project).

Second, we deploy a wide-spectrum network and experimentally evaluate MAWS. Our testbed provides access to the 700 MHz, 900 MHz, 2.4 and 5 GHz bands. These four bands span a spectral range of 5.085 GHz; to the best of our knowledge, this range is the widest to be spanned to date by a single operational access network. We employ our testbed to conduct experiments in both outdoor and indoor environments, under vehicular and pedestrian speeds respectively.

Our experimental evaluation yields the following findings: First, for cross-*spectral* inference, we find that our coupling of propagation models with limited measurements is more tolerant to the fading-induced deviation of measured RSSI when measurements are taken in the most separated bands. Second, for cross-*spatial* inference, we find that RSSI measurements from only two other prior locations suffice for MAWS due to the adequate accuracy of propagation models for the purposes of selecting *diverse* bands. Next, we evaluate MAWS' inference of *spectral usage* for vehicles and pedestrians in two networks with real users. We find that, despite the limited availability of measurements, our method can predict spectral usage with an error of 10-25%. Namely, while packet-scale channel occupancy varies significantly, MAWS' passive second-scale measurements can exploit substantial spatial and temporal usage correlation. We show that despite MAWS' imperfect inferences and selections, it nonetheless attains throughput gains exceeding 100% over scanning, by substituting most scanning time with transmissions. Moreover, compared to the common, near-zero-overhead practice of fixed band prioritization, MAWS enables net throughput gains of up to 120% including overhead.

The rest of this paper is structured as follows. In Sec. II, we present MAWS. We describe our wide-spectrum testbed in Sec. III, and we experimentally evaluate MAWS in Sec. IV. Finally, Sec. V overviews the related work, and Sec. VI concludes this paper.

II. MAWS: Mobile Access of Wide-Spectrum Networks

In this section, we present *MAWS*, a scheme for mobile clients of wide-spectrum networks to evaluate and select their association options, i.e., pairs of APs and channels, which can span multi-band spectrum.

A. MAWS Overview

In MAWS, mobile clients make association and handoff decisions by predicting throughput and delay for each association option. To predict these two metrics, we propose a methodology for mobile clients of wide-spectrum networks to estimate channel quality and spectral usage. Specifically, mobile clients employ a limited-rate scanning process via which they measure channel quality, i.e., link signal strength, in as few as two channels from two different bands. Via this process, clients also sparsely estimate across space each channel's usage, i.e., the fraction of time that other nodes use this channel. Using the limited measurements collected in other bands and at other locations, a MAWS client infers channel quality and spectral usage at its current location for its available bands, despite the lack of a direct measurement. The key technique to infer channel quality is to estimate the cross-spectral correlation of channel quality at each location and the cross-spatial correlation for each band by coupling the sporadic measurements with propagation models. The key technique to infer spectral usage is to exploit the spatial and temporal correlation in the usage of each channel and to account that the spatial correlation increases with decreasing frequency. Finally, our inference-based methodology is coupled with a metric estimating each association option's delay performance, which is also determined by the propagation characteristics of the option's channel. Our inference-based methodology allows MAWS clients to dynamically prioritize APs and multi-band spectrum without exclusively employing scanning or historical values of channel quality and spectral usage. Fig. 1 diagrammatically illustrates the overview of MAWS.

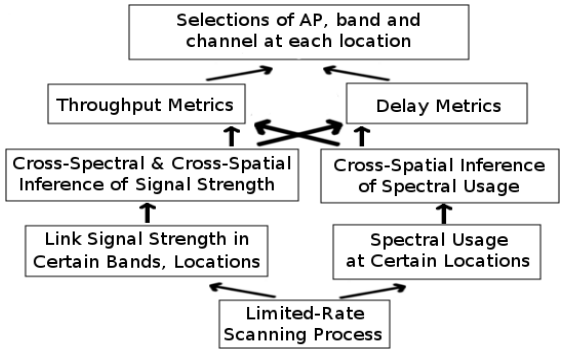

Fig. 1. Diagrammatic Overview of MAWS

B. MAWS Architecture

Our work applies to wide-spectrum access network with mobile clients. We consider that APs and clients are equipped with either a single multi-band radio or multiple single-band radios. Moreover, we consider that all nodes can access a common set of bands and employ the same MAC scheme in all bands. Nevertheless, our work can easily be adapted to the case where the nodes' accessible bands are different; our work can also be extended for networks that do not employ the same MAC scheme in all bands, by incorporating the MAC's impact on throughput and delay performance in the evaluation of association options.

Furthermore, we consider mobile clients to transmit packets via a single radio at a time, even if they are equipped with multiple single-band radios, to conserve energy and increase their lifetime. Moreover, we assume that mobile clients know the coordinates of their locations during their scanning and inference instances. Such localization can be realized via e.g., intermittent activation of the prevailing GPS devices[1] or spectral fingerprinting.[2] In addition, we assume that mobile clients have access to estimates of AP locations obtained from AP-logging databases[3] or announced by APs via their beacons.

[1] 'GPS-Enabled Cell Phones Go Mainstream': http://www.m2mmag.com/
[2] See Place Lab: http://ils.intel-research.net/place-lab
[3] See for example: http://wigle.net/

C. Inferring Channel Quality

In this subsection, we present the two methods that MAWS clients employ for cross-spectral and cross-spatial inference of link signal strength, our considered channel quality metric. Mobile clients employ the cross-spectral inference method at a location where they measure the RSSI of links in two different-band channels and infer signal-strength values for the remaining frequencies. Otherwise, clients employ the cross-spatial inference method to infer the signal strength of a link in a given band, using RSSI measurements collected in the same band at other locations.

Both methods infer the average channel quality of a link at a given location and frequency and do not capture location-dependent and time-varying effects on signal propagation such as shadowing and multipath-fading. Nevertheless, our solution is environment-agnostic and practical, as it does not require a detailed description of the propagation environment, such as terrain maps (see e.g., [6]). At the same time, our employment of measurements can enable more accurate inference than alternatives employing neither information describing the propagation environment nor measurements.

1) Cross-Spectral Inference: Prior measurements and propagation models indicate that signal strength (denoted by P) is inversely proportional to an α-power of frequency f (see, e.g., [7], [8], [9], [10]):

$$P \propto \frac{1}{f^{\alpha}} \tag{1}$$

A frequency-exponent value of $\alpha = 2$ is widely employed for various environments (see, e.g., [7], [8], [9]), while Riback et al. suggest a frequency-dependent selection of α, with $\alpha \in \{2, 2.3, 3\}$ [10]. While an infinite number of functions $P(f)$ can satisfy the relationship specified by Eq. (1) for a given α, only one function $P_l(f)$ represents how signal propagates in each frequency over a link l as a result of the location and composition of obstacles. Specifically, signal strength decreases with frequency more rapidly as obstruction increases. Unfortunately, clients do not necessarily foreknow how the channel quality depends on frequency for every link and every location, and precise modeling of these relationships requires an extensive collection of measurements.

We propose that clients infer channel quality in a given frequency and location by using their limited, same-location RSSI measurements in different bands to characterize individual link models of inverse signal-strength proportionality to an α-power of frequency. To characterize the individual link models, we interpolate, for each link l, the RSSI $P_{l,i}$ measured in frequencies f_i with a function of the form $P(f) = \frac{m}{f^{\alpha}} + b$. This interpolation yields characteristic coefficients m_l and b_l for each link l. Consequently, clients can infer the signal strength of link l in frequencies that are not scanned, using the function: $\hat{P}_l(f) = \frac{m_l}{f^{\alpha}} + b_l$. To enable a low-complexity, analytical solution, we transform the problem to its equivalent linear least squares form by conducting the following variable transformation: $z = \frac{1}{f^{\alpha}}$. Thus, clients can analytically interpolate the RSSI measurements with the function $P(z) = mz + b$ [11].

Our inference method is summarized as follows:

Cross-spectral Signal-strength Inference:

- Consider an input of N channel-quality measurements for the client-AP link l:

  $$(f_i, P_i), \quad i = 1 : N,$$

  collected in at least two different carrier frequencies ($N \geq 2$).
- Consider the transformed variable $z = \frac{1}{f^{\alpha}}$, and the respective measurements: $(z_i, P_i), \quad i = 1 : N$
- Apply the linear least squares method by conducting linear fitting to the points (f_i, P_i) with the function:

  $$\hat{P}_l = m_l \times z + b_l,$$

  where:

  $$m_l = \frac{N \sum_i^N z_i P_i - \sum_i^N z_i \sum_i^N P_i}{N \sum_i^N z_i^2 - (\sum_i^N z_i)^2}$$

  $$b_l = \frac{\sum_i^N P_i - m_l \sum_i^N z_i}{N}$$

- Infer the channel quality of link l in unscanned frequencies f_k:

  $$f_k \neq f_i \quad \forall i = 1 : N,$$

  using the following model for link l:

  $$\hat{P}_l(f_k) = \frac{m_l}{f_k^{\alpha}} + b_l$$

2) Cross-Spatial Inference: Prior measurements and propagation models indicate that signal strength decays logarithmically with the distance d from the transmitter node [7]:

$$P_{dBm}(d) = P_{dBm}(d_0) - 10\gamma \log_{10}\left(\frac{d}{d_0}\right) + \sigma \tag{2}$$

In Eq. (2), $P_{dBm}(d_0)$ is the received signal strength at a reference distance from the transmitter d_0; σ is a zero-mean Gaussian random variable that represents shadowing, i.e., the deviation in P_{dBm} between similar propagation scenarios; finally, γ is the path loss exponent, a parameter representative of the propagation conditions in an environment. This exponent is dependent on frequency and on the location and composition of obstacles [7]. Unfortunately, mobile clients do not necessarily foreknow the path loss exponents γ_f and the extent of shadowing in every frequency f and at every location, and precise estimation of these parameters requires extensive measurement collection.

We propose that mobile clients infer channel quality by utilizing the limited RSSI measurements from previous locations to estimate path loss exponents for each band, and hence to apply the log-distance propagation model. To do so, clients also utilize their location information and estimates of the AP locations. To estimate the path loss exponents, clients apply regression analysis to the points $(d_i, P_{b,i})$ for each band b, where $P_{b,i}$ is the measured RSSI in band b and at a distance d_i from the AP. Our chosen function for

regression analysis expresses the logarithmic decay of signal strength with distance: $P(d) = \beta - \gamma 10 \log_{10}(d)$. We reduce the problem to linear least-squares fitting via the following variable transformation: $z = -10 \log_{10} d$. This process yields coefficients γ_b and β_b approximated analytically for each band b. Consequently, clients can infer the signal strength for a link of length d in band b without scanning using the following function: $\hat{P}_b(d) = \beta_b - 10\gamma_b \log_{10}(d)$. Our inference method is summarized as follows:

Cross-spatial Signal-strength Inference:

- Consider the RSSI P_{ij} that are measured in band b, for AP j, at N locations of prior visit g_i, $\quad i = 1 : N$, all within distance D from the current location g_{cur}: $d(g_{cur}, g_i) \leq D, \quad \forall i$.
- Consider the location G_j of the j'th AP. We denote the Euclidean distance between each location g_i and G_j by d_{ij}.
- Consider the transformation $z = -10 \log_{10} d$
- Mobile clients estimate the path loss exponent $\gamma_{b,j}$ for band b and AP j for the neighborhood within distance D from the current location:
  Apply the linear least squares method by conducting linear fitting to the points:

$$(z_{ij}, P_{ij}) = (-10 \log_{10} d_{ij}, P_{ij}), \quad i = 1 : N,$$

  with the function:

$$P(z) = \gamma_{b,j} \times z + \beta_{b,j},$$

  where:

$$\gamma_{b,j} = \frac{N \sum_i^N z_{ij} P_{ij} - \sum_i^N z_{ij} \sum_i^N P_{ij}}{N \sum_i^N z_{ij}^2 - (\sum_i^N z_{ij})^2}$$

$$\beta_{b,j} = \frac{\sum_i^N P_{ij} - \gamma_{b,j} \sum_i^N z_{ij}}{N}$$

- Infer the channel quality of the link to AP j at location g_i using the approximated log-distance propagation model:

$$\hat{P}(d_{ij}) = \beta_{b,j} - 10 \times \gamma_{b,j} \times \log_{10}(d_{ij})$$

- Repeat the method for all bands and all APs.

D. Inferring Spectral Usage

Here, we propose a method for mobile clients to infer, at a given location, each channel's usage, i.e., the fraction of time that the channel is used by other nodes.

Exploiting cross-spatial and -temporal correlation: Our method exploits the correlation in sensing a signal at neighboring locations and the temporal correlation in the usage of a channel. Towards this end, clients infer spectral usage from a limited number of usage estimates that they collect at locations of prior visit. Specifically, the usage of the channel centered at frequency f is inferred at current location g_c as a weighted average of that channel's measured usage at locations g_i:

$$\hat{u}_f(g_c) = \frac{\sum_{i=1}^N u_f(g_i) * w_f(g_i, g_c)}{\sum_{i=1}^N w_f(g_i, g_c)} \tag{3}$$

Usage estimates are denoted by $u \in [0, 1]$, with greater u values denoting higher usage.

In Eq. (3), individual weights are assigned to each usage estimate, as different pairs of locations exhibit dissimilar spatial correlation in the usage of a channel. For instance, the probability that two locations share a common interferer decreases with the distance of the locations. In addition, this probability decreases with frequency, as the interference range also decreases. A weight assignment that captures the abovementioned relationships is the following:

$$w_f(g_i, g_c) = \max\{I_b - d(g_c, g_i), 0\}, \tag{4}$$

where $d(g_i, g_j)$ is the Euclidean distance between two locations g_i, g_j; I_b is an estimate for the interference range in band b and its estimation follows.

Estimating Interference Range: Multiple factors affect the interference range of a client. In wide-spectrum networks, selecting a lower-frequency band can yield a dramatic increase in interference range. Of course, the interference range also depends on the location of the interferer and the client, as distinct propagation conditions are determined by the obstruction inbetween each pair of locations. Finally, the interference range may also vary over time, as a result of channel fading. Unfortunately, clients do not necessarily know the location of their interferers in each band or their interference range at each location, time and frequency.

We propose that mobile clients of wide-spectrum networks approximate a single interference range for each band by employing the cross-spatial method for channel-quality inference (see Sec II-C2). Specifically, clients utilize RSSI measurements to infer the channel quality of client-AP links at different locations. Thus, clients can estimate the maximum distance I_{bj} from AP j at which signals are received at interfering power levels, in band b:

$$I_{bj} = \max_d \left\{ \hat{P}_{bj}(d) > P_{int} \right\}$$

We denote by $\bar{I}_b$ the average range of AP-generated interference I_{bj}, averaged over all APs. In our scheme, clients consider $\bar{I}_b$ as the estimate of their interference range in band b.

E. Limited-Rate Scanning

MAWS clients employ limited-rate scanning to sparsely measure the highly variable metrics of channel quality and available airtime across spectrum and space. Nonetheless, these coarse-grained samples suffice to enable, via our methods, inference of these two metrics for a wide spectral and spatial range. In general, MAWS is a framework encompassing a broad set of scanning parameters, which can be adjusted to client velocity and spectrum availability.

Channel Quality: Clients periodically measure channel quality by probing every π_{cq} seconds. At each probing action, clients select two bands to probe and probe any single channel of each selected band. Probing two bands minimizes the scanning required to enable our cross-spectral inference of channel quality. Moreover, different pairs of bands are selected over

time to collect measurements enabling cross-spatial inference of channel quality for all bands. By receiving probe responses, clients measure the RSSI $P_{k,f}(g_i)$ for their links to each AP k at locations g_i and frequency f.

Spectral usage: Additionally, clients passively estimate spectral usage via periodical sniffing. Every π_u seconds, clients sniff a single channel; each sniffing action lasts t_{snif} s., and channels are selected sequentially. As a result, clients calculate usage estimates $u_f \in [0,1]$ for the respective frequency f, where u_f equals the fraction of the sniffing duration that clients sense the channel as used by other nodes.

F. Selection of Association Options

Mobile clients may individually prioritize throughput versus delay performance, as they may dissimilarly tolerate packet delay, which includes handoff, transmission and contention delay. Denote δ as the delay sensitivity of a client, $\delta \in [0,1]$, such that delay tolerance decreases with δ.

Given the inferred and measured metrics of channel quality and spectral usage, MAWS clients estimate throughput and delay metrics for their numerous association options at each location, to select the one that best meets their performance objectives.

Throughput Metric: For association with AP k in frequency f, throughput is predicted via estimates for the attainable link rate and the available airtime of that frequency:

$$T_{k,f}(g_t) = R_i\left(P_{k,f}(g_t), W_i\right) \times \left[1 - u_f(g_t)\right], \qquad (5)$$

In Eq. (5), W_i is the channel width of the band i that includes frequency f, and R_i is the attainable rate under interference-free conditions. The dependence of R_i on channel width and channel quality metrics, such as signal strength, can be empirically estimated. Finally, $P_{k,f}(g_t)$ and $u_f(g_t)$ are the inferred or measured metrics of channel quality and usage for location g_t, respectively.

Delay Metric: Delay under each association option is estimated via a metric incorporating determinant factors of delay performance in a wide-spectrum network. Denote $D_{k,f} \in [0,1]$ as the delay metric for the client's association to AP k in frequency f.

First, the metric incorporates the delay incurred by the selection of an association option; selecting a different band for the currently associated AP incurs a significantly shorter delay than handing off to a different AP. Specifically, the channel switching delay is orders of magnitude lesser than handoff delay (e.g., 80 μs vs. 25-800 ms; values reported in [12], [2], [3]). The key reason for this difference is that handoffs in widely employed MAC schemes such as 802.11a/b/g typically require association handshakes.

Moreover, the delay metric penalizes association options increasingly with frequency, as the handoff rate increases with decreasing coverage. Finally, the delay metric employs inferred or measured metrics of channel quality and spectral usage to estimate the transmission and access delay (e.g., contention delay) under each association option. We provide the complete presentation of the delay metric in the Appendix, Sec. C.

Selection of Association Options: The throughput and delay metrics are weighted according to the client's delay sensitivity to yield a joint metric that drives the selection of association options and expresses the relevance of each option to the client's performance objectives:

$$j_{k,f}(t) = T_{k,f}(g_t) \times \left[1 - \delta \times D_{k,f}(g_t)\right], \qquad (6)$$

Finally, clients use a hysteresis threshold η to refrain from invoking handoffs to candidate associations (k,f) with marginally higher joint metrics $j_{k,f}$ than that of their current association (AP_{cur}, f_{cur}). Hence, clients select a different association (k,f), when:

$$j_{k,f} > j_{AP_{cur},f_{cur}} + \eta$$

Selections are followed by a handoff to another AP when $k \neq AP_{cur}$ and by a channel switching when $k = AP_{cur}$ and $f \neq f_{cur}$.

G. Example Access Model

Similarly to many network architectures in which nodes utilize more channels than the number of their radios (e.g. [13]), nodes of wide-spectrum networks can coordinate through a control channel. Specifically, two nodes can first exchange packets in the control channel (such as RTS, CTS, probe requests); such packets specify another channel, in which the nodes switch and they further exchange packets (such as data packets, probe responses); subsequently, they switch back to the control channel.

This access model can be realized by a dedicated control channel (e.g., [13]) or by dynamic control channels, which can be established with channel-hopping protocols (e.g., [14]). While bandwidth reduction has been a counter-argument to control channels in single-band networks, spectral resources abound in wide-spectrum networks.

III. Experimental Platform

This section describes the wide-spectrum network that we deploy and utilize to experimentally evaluate MAWS. Our testbed provides access to four spectral bands: the 700 MHz band,[4] and the ISM bands of 900 MHz, 2.4 and 5 GHz. These bands span a spectral range of 5.085 GHz; to the best of our knowledge, this range is the widest to be spanned to date by a single operational access network.

To evaluate MAWS, we deploy a two-AP network accessed by a mobile and a static node. The mobile node is either moving in a car at vehicular velocity or is placed on a cart while moving at pedestrian speeds, depending on the experiment. The other client is stationary and serves the purpose of injecting traffic into the network. APs are placed approximately 15 meters above the ground at two different balconies of Duncan Hall, at Rice University in Houston, TX.

[4]The 700 MHz band is allocated for public-safety services. However, it was not used at the deployment area.

In our testbed, nodes are equipped with multiple single-band radios, one for each of the four bands to which access is provided. Fig. 2 illustrates a MAWS node, used for both APs and clients. Finally, the client nodes employ low-gain, 5 dBi

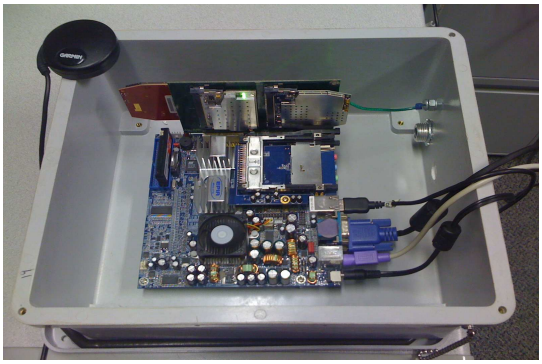

Fig. 2. A 4-radio MAWS node

antennas for all bands, while APs employ 9 dBi antennas to provide a wider coverage.

Testbed specifications: Our platform is x86-based, using Gentoo Linux kernel 2.6.34. Nodes are equipped with Atheros chipset mini-PCI interfaces from the Ubiquiti Networks XtremeRange series. These are the XR7, XR9, XR2 and XR5 radios, which operate in 760-780 MHz, 902-927 MHz, 2.401-2.483 GHz and 5.160-5.845 GHz, respectively. The XR2, XR7 and XR9 interfaces use a 802.11g MAC, while XR5 follows the 802.11a standard. The interfaces function with the `ath5k` open-source driver.

IV. Evaluation of MAWS

In this section, we evaluate MAWS using the deployment described in Sec. III. We study the individual components of MAWS as well as their joint interaction. Thus, we assess the accuracy of the methods inferring channel quality and spectral usage, and we compare MAWS with alternatives for mobile access, such as scanning and fixed-band prioritization.

A. Inferring Channel Quality

To assess the accuracy of the cross-spectral and cross-spatial methods, we design the following experiment. First, we deploy a single-AP wide-spectrum network, in which a mobile client collects RSSI measurements of its link to the AP; the client collects a *single* RSSI measurement at each of multiple locations and for each of the four bands of our testbed. Next, we provide only a subset of the collected measurements as an input to the two inference methods. Then, the methods infer RSSI values for the frequencies and locations that are not included in the subset. These inferred values are compared with the respective RSSI measurements. Specifically, we place the single AP on a third-floor balcony of Duncan Hall, in Rice University. The mobile client collects RSSI measurements at 773 MHz, 912 MHz, 2.447 GHz and 5.2 GHz. This experiment is conducted twice, once outdoors with a car as a mobile client, and once indoors with the client placed on a cart. In both cases, the selected measurement locations yield client-AP links that cover a wide-range of factors affecting signal propagation such as distance and intermediate obstruction. Fig. 3 depicts the deployment and the selected locations. We collect all outdoor

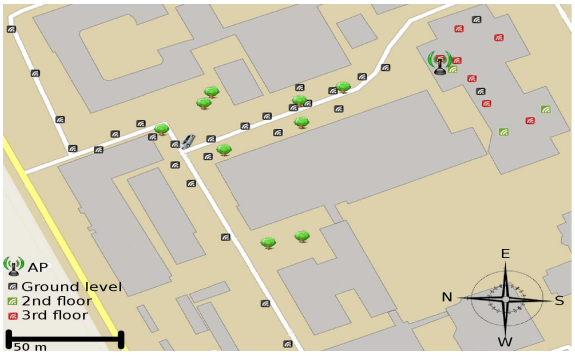

Fig. 3. Measurement Locations

measurements in the same day and all indoor in the next day; successive measurements are taken at least 15 minutes apart.

1) Cross-Spectral Inference: During each scanning action, MAWS clients probe two channels belonging to different bands (see Sec. II-E). Here, we assess how accurately clients can infer channel quality for the remaining, non-scanned bands and how this accuracy can increase under appropriate selection of the probed channels. Thus, the input of our inference method consists of two same-location RSSI, measured in two of the four frequencies considered in the experiment; for the same location, the inferred RSSI values for the remaining two frequencies are compared with the respective measurements. We repeat the experiment for each location and for every possible combination of input selection, i.e., for every choice of two channels out of the four in which measurements are collected. To distinguish between different input scenarios, we associate each combination with a measure of dissimilarity between the propagation characteristics of the two chosen frequencies f_1, f_2. This measure is driven by models suggesting the inverse proportionality of channel quality to an α-power of carrier frequency (Eq. 1), and it is given by $|\frac{1}{f_1^\alpha} - \frac{1}{f_2^\alpha}|$.

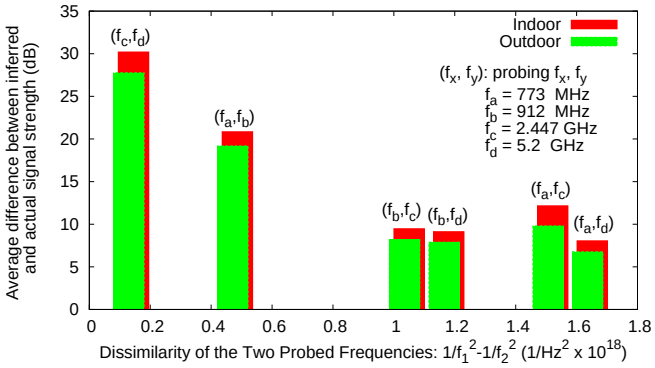

Fig. 4. Impact of Scanning Frequencies on Inference of Channel Quality

We calculate the difference between the dBm values of actual and inferred RSSI as well as the absolute value of the difference. Fig. 4 depicts the average such absolute value, averaged over all instances of inference, as a function of the dissimilarity of the two probed frequencies. The results shown in Fig. 4 are obtained for $\alpha = 2$ (see Eq. (1)). Fig. 4

indicates that the accuracy of our interpolation-based inference method increases with the degree of dissimilarity of the two probed frequencies. Specifically, we show in the Appendix, Sec. A, that our method infers the time-average channel quality of a link with an error that increases with $\frac{\Delta_2 - \Delta_1}{\frac{1}{f_2^\alpha} - \frac{1}{f_1^\alpha}}$, where Δ_i is the difference between an RSSI measurement and the time-average channel quality in frequency f_i. Thus, the tolerance of our method to fading-induced deviation of the RSSI measurements from their average value increases with the dissimilarity of the two probed frequencies.

Finally, Fig. 4 illustrates that under an appropriate selection of the probed frequencies, our method can infer the RSSI values in other bands within 7 dB of the actual RSSI. To add perspective to this difference, we empirically compare the link rates attainable under two RSSI values that differ by 7 dB (see Appendix, Sec. B). We find that the difference of the rates is upper bounded by 17% of the highest attainable rate; moreover, the compared rates are equal under half of the possible values for two RSSI differing by 7 dB. While a non-negligible error, such inferences nonetheless suffice for MAWS' selection objectives without sacrificing additional airtime in scanning (see Sec. IV-C). *Finding: Using only two RSSI measurements in two different bands and an appropriate selection of probed frequencies, MAWS clients can infer same-location channel quality in other bands within 7 dB of the actual dBm value.*

2) Cross-Spatial Inference: MAWS clients can infer the channel quality in a given band and location using same-band RSSI measurements from other locations (Sec. II-C1). Here, we assess the error of this inference as a function of the number of RSSI measurements and the robustness of our method to inaccurate knowledge of the AP locations. Towards this end, we first consider all possible subsets N_{in} of the N locations depicted in Fig. 3, $|N_{in}| \in [2, N-1]$. For each band, our inference method employs the single RSSI measurement at each of the N_{in} locations, estimates a distinct path loss exponent and infers channel quality at the remaining $N - N_{in}$ locations. The inferred values are compared with the respective RSSI measurements, which are not employed by our method. We apply our methodology separately for the two different environment scenarios, i.e., indoor and outdoor. Moreover, we repeat the experiment under different values of ϵ, a parameter representing the difference between the assumed and the actual length of the client-AP link.

We calculate the difference between the dBm values of actual and inferred RSSI, and the absolute value of the difference. Fig. 5 depicts the average such value in inferring channel quality at the $N - N_{in}$ locations as a function of the number of RSSI measurements N_{in}. The figure depicts results from cross-spatial inferences in one band (900 MHz), while we observe similar trends and findings for all other bands. As Fig. 5 indicates, two RSSI measurements at different locations enable inference of channel quality within $0-6$ dB of the value inferred via RSSI measurements in $N-1$ locations. Moreover, our method's accuracy increases with link length

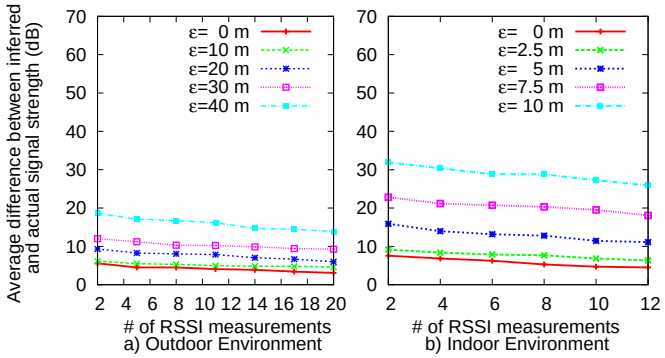

Fig. 5. Impact of RSSI Availability on Inference of Channel Quality

(longer links in the outdoor case) as signal strength decreases logarithmically with distance; thus, for a given ϵ, the difference $P(d+\epsilon) - P(d)$ decreases with d. *Finding: Inference of channel quality with two RSSI measurements is marginally less accurate than inference with the highest number of measurements.*

B. Inferring Spectral Usage

To assess the accuracy of MAWS in inferring spectral usage, we design the following experiment. A MAWS node traverses the coverage area of a network providing access to real users at channel c. Using *kismet*, the node continuously sniffs this single channel c. We discretize time in seconds; with each second t_i, we associate a single location g_i chosen among the possibly many visited during t_i. For each of the second-long time-intervals t_i, we calculate the fraction of time $u_f(g_i)$ that the center frequency f of channel c is sniffed as busy, i.e., the fraction that *kismet* reports packet transmissions in c. Then, we consider a subset of the entire sniffing process, a subset representing periodic sniffing instances, which last t_{snif} s. and repeat every π_{ut} s. These sniffing instances comprise the input of our inference method (Sec. II-D), which yields usage inferences $\hat{u}_f(g_i)$ for f, at each location g_i. We compare the inferred usage $\hat{u}_f(g_i)$ with the actual $u_f(g_i)$, for each i. We generalize our single-frequency assessment for the case of a wide-spectrum network with N channels, considering a sniffing pattern where a client sniffs a different channel every π_{ut} seconds and the same channel every $N \times \pi_{ut}$ seconds.

Networks with real users: To evaluate our inference method, we conduct experiments in operational networks used by real clients. We consider two networks that provide access to two different bands: *(i)* The TFA network, an operational, urban mesh network consisting of approximately 20 APs and providing access to channel 11 of the 2.4 GHz band. During the experiments, clients maintain vehicular speeds while sniffing this channel, to traverse the network coverage along a 4.1 km route within 800 s. *(ii)* A university network providing indoor access to the 2.4 and 5 GHz bands in Duncan Hall, Rice University. We conduct the experiment once for each band, sniffing channel 6 in the 2.4 GHz band and channel 48 in the 5 GHz band. In these two experiments, the client traverses the network coverage at pedestrian speeds along a

200 m route, within 260 s. All experiments are conducted during hours that actual users access the networks.

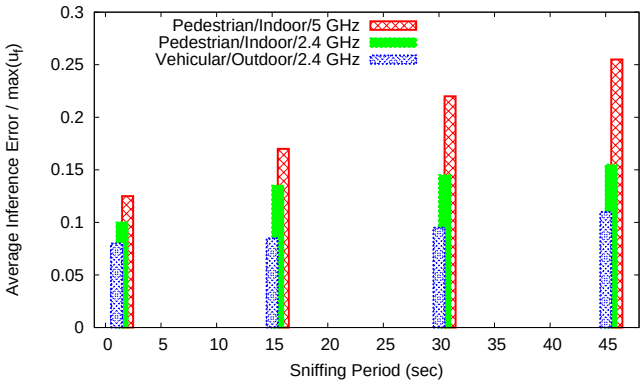

Fig. 6. Impact of Sniffing Rate on Usage Inference ($t_{snif} = 100$ ms)

For each of the three velocity/band scenarios, we calculate the inference error $\epsilon(g_i) = |\hat{u}_f(g_i) - u_f(g_i)|$ at each location g_i. Moreover, we calculate the average inference error $\bar{\epsilon}$, averaged over all locations g_i, as a fraction of the maximum usage $u_{max} = \max\{u_f(g_i)\}$ that is measured during each experiment. Fig. 6 depicts this normalized metric of average error, as a function of the sniffing period, for $t_{snif} = 100$ ms. Under realistic traffic patterns, the usage of a frequency is strongly correlated across space and time. Thus, despite the existing variations in the usage of a frequency, the average inference error $\bar{\epsilon}$ of our method ranges within 10-25% of the maximum measured usage, even under a very infrequent sniffing period, e.g. $\pi = 45$ s. We consider that the usage of N channels can be inferred with such an accuracy by periodically and sequentially sniffing the channels every $\frac{\pi}{N}$ s (our wide-spectrum testbed can operate in up to 40 channels). *Finding: Under realistic traffic patterns, mobile clients can infer spectral usage along their trajectory with an average error of 10-25% of the maximum measured value.*

C. MAWS vs. Alternatives for Mobile Access

Here, we compare MAWS with two alternatives for mobile access: *(i)* exclusive employment of scanning, and *(ii)* fixed-band prioritization. Association based on scanning alone can result in more accurate estimates of channel quality and spectral usage than our inference-based methodology, which is subject to inference errors. However, this accuracy comes at the cost of sacrificing airtime that increases with the number of scanned channels. Alternatively, clients can omit scanning and employ static band prioritizations, similarly to today's common practice for selecting between different networks operating in diverse bands. However, such an approach is oblivious to the underlying usage and channel quality of each band at any given location.

Towards these two comparisons, we design the following experiment. We place two MAWS APs at different locations of the university campus, and a mobile client experimentally measures its throughput performance via a UDP *iperf* session. Specifically, we repeat the measurements along the same trajectory for each possible selection of AP and spectral band (eight combinations for two APs and four bands). Moreover, we repeat the experiment under different usage values for each band; thus, we control band usage by conducting the experiment at hours that no actual users access the medium and by injecting controlled traffic via another, static client. In all scenarios, the client can passively measure RSSI from its associated AP via the exchanged data packets. Finally, we repeat the experiment for different durations of periodic inactivity, i.e., pausing of the *iperf* session; such idle intervals represent the airtime sacrificed to estimate channel quality and usage in other channels.

As our platform is limited to only 2 APs and 2 clients, we utilize the collected measurements to emulate client performance under each mobile-access scheme in larger-network scenarios. For our emulation, we consider a linear trajectory and a placement of 10 APs. Each AP is placed in such a manner that its relative position to a certain segment of the trajectory is representative of the experiment conducted for measuring throughput performance. Moreover, the AP placement enables multiple scenarios of coverage overlapping. In our emulation model, many different association options exist in each location of the trajectory; our experiments have measured the channel quality of each option and the offered throughput under many scenarios of spectral usage and scanning frequency. Our emulation model assigns to each location the AP-channel pair chosen by each mobile-access scheme and associates the respective empirical throughput.

1) MAWS vs. Scanning: As MAWS has errors in inferring channel quality and spectral usage, we perturb the actual measured and controlled values of these metrics according to the inference errors quantified in Sec. IV-A and Sec. IV-B. In contrast, we consider that the scanning-only alternative has error-free measurements of channel quality and usage. Both schemes select their own association options at each location according to Eq. (6), for $\delta = 0$. Then, the MAWS throughput at each location is given by the one measured under the selected association during periodic inactivity for $t_{snif} + 2 \times t_{prob}$ s., as MAWS periodically sniffs one and probes two channels. Moreover, scanning's throughput is given by the one measured during periodic inactivity of $N \times t_{prob}$, where N is the number of channels in the network.

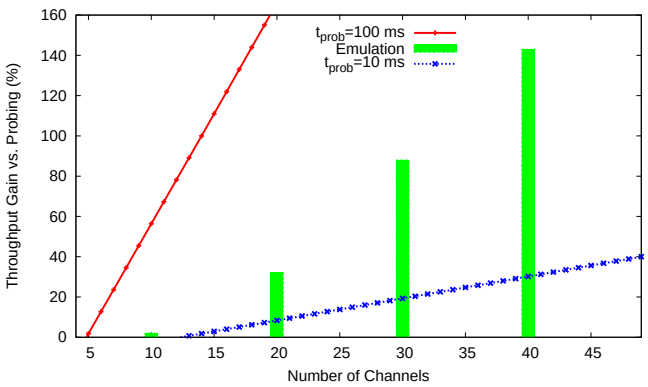

Fig. 7. MAWS vs. Exclusively Scanning

Fig. 7 depicts the throughput gain of MAWS as a function

of the number of channels, for $t_{prob} = 25$ ms, $t_{snif} = 100$ ms and a scanning period of 1 second. To add perspective, it also depicts an analytically calculated gain for different values of t_{prob}. We find that the airtime sacrificed by scanning overwhelms its advantage against MAWS, i.e., the more accurate evaluation of association options, when the network operates in more channels than single-band networks (e.g., 11). In multi-band networks, the MAWS gain can reach up to 140% for 40 channels; specifically, it increases linearly with the number of channels, as so does the airtime consumed by scanning, while the accuracy of our inference methods remains the same (see Sec. IV-A, IV-B). *Finding: Throughput gains over exclusive employment of scanning can exceed 100%.*

2) MAWS vs. Fixed-Band Prioritization: We consider three versions of fixed-band prioritization: selection of the highest-frequency band (among those available), selection of the lowest-frequency, and highest preference of the 2.4 GHz band with arbitrary preference order for the remaining bands. When multiple APs provide coverage at the same location and in the most preferred available band, we consider that these three policies select the highest-RSSI AP. We consider the 10-AP network scenario for many different cases of spectral usage, which is controlled in our experimental methodology. We assign the empirical throughput performance to MAWS clients as described; moreover, we assign to each fixed-band prioritization policy the respective throughput values, measured under no periods of inactivity (no scanning cost).

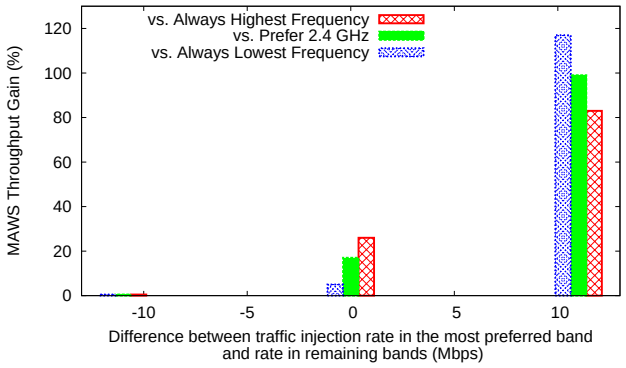

Fig. 8. MAWS vs. Fixed Band Prioritization

Fig. 8 depicts the throughput gain of MAWS over each of the fixed-prioritization policies. The x-axis represents the difference between the rate of injected traffic in the highest-preference band R_{pref} and the rate for all remaining bands R_{rest}. *Finding: Despite scanning, MAWS yields significantly higher throughput performance than fixed-band prioritization when the usage of the statically preferred band is equal or higher than that of the remaining bands.* The reason is that MAWS incorporates estimates of spectral usage and channel quality in the selection of association options, despite those estimates being imperfect. Specifically, the MAWS gain due to the incorporation of spectral-usage estimates increases linearly with the additional traffic rate in the statically preferred band; for an additional rate of 11 Mbps, the gain can reach up to 120%. Under identical band usage, MAWS incurs gain due to the incorporation of channel-quality estimates; such a gain increases with the frequency of the most preferred band and can reach up to 25%.

V. Related Work

Related work can be classified into: *(i)* mobile access of single-band networks, and *(ii)* non-mobile access of single- and multi-band wireless networks.

A. Mobile Access

IEEE Standards: The 802.21 standard enables handoff between different network technologies [15]. In addition, 802.11p is a standard for vehicular communication in the dedicated spectral band of 5.9 GHz [16].

Overlaid Cells: In cellular networks, multiple overlaid layers of coverage are provided by cells of different sizes that operate in the same band. In such networks, cell selection is driven by client velocity (e.g., see [17]). At the same time, mobile clients widely employ fixed-band prioritization to select between independent networks operating in dissimilar bands (e.g., smart-phone preference of W-LAN to 3G; see also [1]).

Non-Cellular Single-Band Networks: In our prior work, vehicular clients of multi-hop wireless networks prolong associations to better performing APs by accounting for disparities in offered AP throughput [3]. Deshpande et al. propose the disengagement of mobile clients from scanning via the utilization of historical RSSI values at revisited locations [4]. In cognitive networks that operate in channels exhibiting identical propagation characteristics, mobile secondary users opportunistically access spectrum that is not occupied by primary users [18] and enable collaborative spectrum sensing [19].

In contrast to prior work in mobile access, MAWS dynamically prioritizes diverse spectrum, without exclusively employing historical data or scanning.

B. Non-mobile Access

Diverse-Spectrum Networks: Shu et al. address throughput maximization by regulating the scanning frequency of cognitive radios [20]. Moreover, prior work addresses spectrum access in networks operating specifically in the UHF white-spaces (see, e.g., [21]). To predict channel availability, Chen et al. employ historical measurements ranging from 20 MHz to 3 GHz and identify patterns of channel usage [5], while Tumuluru et al. employ methodologies based on neural networks and Markov chains [22].

Single-Band Networks: Deployed urban-scale networks can be evaluated via propagation models, terrain maps and measurements obtained within a wide range of locations [6].

In contrast to prior work in non-mobile access, MAWS considers client mobility; clients are temporarily present at a given location and evaluate their association options via cross-spatial and -spectral inference methods employing measurements collected at other locations. In contrast to all prior work, MAWS is the first scheme designed for evaluation and selection of APs and channels by mobile clients of wide-spectrum networks.

VI. Conclusions

In this paper, we present MAWS, the first scheme designed for mobile clients to evaluate and select their association options in wide-spectrum networks. The key technique in MAWS is for clients to infer channel quality and spectral usage for their current location and bands using limited measurements collected in other bands and at other locations. We experimentally evaluate MAWS using a four-band wide-spectrum network that we deploy. Our evaluation reveals that MAWS yields significant throughput gains over alternatives for mobile access.

References

[1] M. Stemm and R. Katz, "Vertical handoffs in wireless overlay networks," *Mobile Networks and Applications*, vol. 3, no. 4, pp. 335–350, 1998.

[2] V. Mhatre and K. Papagiannaki, "Using smart triggers for improved user performance in 802.11 wireless networks," in *Proc. ACM MobiSys*, Uppsala, Sweden, June 2006.

[3] A. Giannoulis, M. Fiore, and E. W. Knightly, "Supporting vehicular mobility in urban multi-hop wireless networks," in *Proc. ACM MobiSys*, Breckenridge, CO, June 2008.

[4] P. Deshpande, A. Kashyap, C. Sung, and S. Das, "Predictive methods for improved vehicular WiFi access," in *Proc. ACM MobiSys*, Kraków, Poland, June 2009.

[5] S. Yin, D. Chen, Q. Zhang, M. Liu, and S. Li, "Mining spectrum usage data: A large-scale spectrum measurement study," *IEEE Transactions on Mobile Computing*, vol. 11, no. 6, pp. 1033–1046, June 2012.

[6] J. Robinson, R. Swaminathan, and E. W. Knightly, "Assessment of urban-scale wireless networks with a small number of measurements," in *Proc. ACM MobiCom*, San Francisco, CA, Sept. 2008.

[7] T. S. Rappaport, *Wireless Communications. Principles and Practice*, 2nd ed. Prentice Hall, 2001.

[8] Propagation data and prediction methods for the planning of indoor radiocomm. systems and radio local area networks in the frequency range 900 MHz to 100 GHz. Recommendation ITU-R P.1238-7, 2012.

[9] Propagation data and prediction methods for the planning of short-range outdoor radiocomm. systems and radio local area networks in the frequency range 300 MHz to 100 GHz. Recommendation ITU-R P.1411-6, 2012.

[10] M. Riback, J. Medbo, J.-E. Berg, F. Harrysson, and H. Asplund, "Carrier frequency effects on path loss," in *Proc. IEEE Vehicular Technology Conference (VTC)*, Melbourne, Australia, May 2006.

[11] S. Chatterjee and A. Hadi, *Regression Analysis by Example*, 4th ed. Wiley, 2006.

[12] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: slotted seeded channel hopping for capacity improvement in IEEE 802.11 ad-hoc wireless networks," in *Proc. ACM MobiCom*, New York, NY, Sept. 2004.

[13] S. Wu, C. Lin, Y. Tseng, and J. Sheu, "A new multi-channel MAC protocol with on-demand channel assignment for multi-hop mobile ad hoc networks," in *Proc. ISPAN*, Dallas, TX, Dec. 2000.

[14] K. Bian, J.-M. Park, and R. Chen, "Control channel establishment in cognitive radio networks using channel hopping," *IEEE JSAC*, vol. 29, no. 4, pp. 689–703, Apr. 2011.

[15] "IEEE standard for local and metropolitan area networks- part 21: Media independent handover," *IEEE Std 802.21-2008*, pp. c1 –301, Jan. 2009.

[16] *IEEE Std 802.11p-2010: Wireless Access in Vehicular Environments*, pp. 1 –51, July 2010.

[17] X. Wu, B. Murherjee, and D. Ghosal, "Hierarchical architectures in the third-generation cellular network," *IEEE Wireless Communications*, vol. 11, no. 3, pp. 62 – 71, June 2004.

[18] A. Min, K. Kim, J. Singh, and K. Shin, "Opportunistic spectrum access for mobile cognitive radios." in *Proc. IEEE INFOCOM*, Shanghai, China, Apr. 2011.

[19] S. Jana, K. Zeng, and P. Mohapatra, "Trusted collaborative spectrum sensing for mobile cognitive radio networks," in *Proc. IEEE INFOCOM*, Orlando, FL, Mar. 2012.

[20] T. Shu and M. Krunz, "Throughput-efficient sequential channel sensing and probing in cognitive radio networks under sensing errors," in *Proc. ACM MobiCom*, Beijing, China, Sept. 2009.

[21] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with Wi-Fi like connectivity," in *Proc. ACM SIGCOMM*, Barcelona, Spain, Aug. 2009.

[22] V. K. Tumuluru, P. Wang, and D. Niyato, "Channel status prediction for cognitive radio networks," *Wireless Communications and Mobile Computing*, vol. 12, no. 10, pp. 862–874, July 2012.

[23] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proc. ACM MobiSys*, Uppsala, Sweden, June 2006.

Appendix

A. On the Error of Cross-spectral Inference

Lemma 1: Consider a link l and let P_1, P_2 denote two RSSI measured for l in frequencies f_1 and f_2, respectively; also consider the transformation $z = \frac{1}{f^\alpha}$. Our cross-spectral inference method yields a link model $\hat{P}_l(z)$ whose slope differs from that of the time average channel quality $P_l(z)$ by $\frac{\Delta_2 - \Delta_1}{z_2 - z_1}$ $(= \frac{\Delta_2 - \Delta_1}{\frac{1}{f_2^\alpha} - \frac{1}{f_1^\alpha}})$, where Δ_i is the difference between P_i and its time-average value.

Proof: Consider that m, b exist to satisfy Eq. (1), i.e., $P_l(z) = m \times z + b$. Our cross-spectral inference method (Sec. II-C1) will determine a link model $\hat{P}_l = \hat{m} \times z + \hat{b}$, where

$$\hat{m} = \frac{P_2 - P_1}{z_2 - z_1} = \frac{P_l(z_2) + \Delta_2 - P_l(z_1) - \Delta_1}{z_2 - z_1} =$$

$$\frac{m \times z_2 + b + \Delta_2 - m \times z_1 - b - \Delta_1}{z_2 - z_1} =$$

$$\frac{m \times (z_2 - z_1) + \Delta_2 - \Delta_1}{z_2 - z_1} = m + \frac{\Delta_2 - \Delta_1}{z_2 - z_1}.$$

Thus, $\hat{m} - m = \frac{\Delta_2 - \Delta_1}{z_2 - z_1} = \frac{\Delta_2 - \Delta_1}{\frac{1}{f_2^\alpha} - \frac{1}{f_1^\alpha}}$.

■

B. Comparison of Actual and Estimated Link Rates

Claim 1: The difference between the estimated attainable link rate under actual and inferred RSSI values is upper bounded by 17% of the highest attainable link rate. Moroever, for half of the cases in the [-95, 0] dBm range, the inferred and actual RSSI values yield identical estimates of attainable link rate.

Demonstration: To compare the link rates $R(s_1)$ and $R(s_2)$ attainable under the inferred (s_1) and actual channel quality (s_2) we design the following experiment. First, we empirically estimate $R(.)$, by measuring link rate under many values of channel quality. We connect a sender and a receiver via a coaxial cable to enable packet transmissions under interference-free conditions. To measure link rate under different channel quality conditions, we gradually increase the attenuation of the link by adding cascaded attenuators to it. For each of the different attenuation scenarios, we extensively measure UDP throughput using *iperf*.[5] We repeat the experiment for every available transmission rate, and we record the highest attained link rate. We validate this empirical

[5]http://iperf.sourceforge.net/

estimate for all four bands in which our testbed operates by repeating the measurement of UDP throughput in each of the bands.

Our measurements lead to conclusions that are consistent with the literature (e.g., [23]): $R(.)$ can be approximated as a piecewise linear function, where one region corresponds to zero throughput, another to the maximum throughput, and in an intermediate region, UDP throughput can be approximated as a linear function of channel quality in dBm. Specifically, our empirical estimate of $R(s)$ is the following:

$$R(s) = \begin{cases} 0 \text{ Mbps}, & s < -75 \text{ dBm} \\ \frac{7}{8} \times s + 65.6 \text{ Mbps} & -75 \leq s \leq -35 \text{ dBm} \\ 35 \text{ Mbps} & s > -35 \text{ Mbps dBm} \end{cases}$$

Consider 1-dBm increments of s from -95 dBm to 0 dBm. Then for 48% of the cases, $R(s+7)$ equals $R(s)$. For the rest of the cases, $R(s+7) - R(s) \leq 6$ Mbps. 6 Mbps is 17% of the highest attainable rate, which is 35 Mbps. ■

C. Delay metric

Here, we propose a metric indicating the delay performance of a mobile client under each association option in a wide-spectrum network, i.e., under each pair of an AP and a carrier frequency.

Multiple factors can affect the delay experienced by the packets of a client: the handoff frequency and duration, as well as the transmission and access delay, i.e., the time to successfully transmit a packet. In wide-spectrum networks, clients can switch their association with an AP to a different frequency at a significantly smaller delay cost than initiating a handoff to a different AP. Furthermore, carrier frequency affects the tendency of a client to handoff, as well as the transmission and access delay. Specifically, as frequency increases, the number of handoffs initiated by a client also increases, because AP coverage decreases. Moreover, the transmission delay increases with frequency, as channel quality decreases with the latter. Finally, access delay may decrease with frequency, as so does the interference range of the client.

We propose a metric for each association option that incorporates the abovementioned factors affecting delay performance; this metric expresses how carrier frequencies of dissimilar propagation characteristics may affect the delay performance of a client. Specifically, clients calculate at each location g_i the delay-driven metric $D_{k,f}$ for each pair of an AP k and a carrier frequency f, which is given by:

$$D_{k,f}(g_i) = \frac{H(k,f) + t_h(f) + \hat{u}_f(g_i) + d_{tr}(\hat{P}_{k,f}(g_i))}{4} \quad (7)$$

In Eq. (7), all variables range in $[0, 1]$, and their higher values indicate higher delay. As a metric indicative of the access delay in frequency f and at location g_i, we consider the respective usage metric $\hat{u}_f(g_i)$, which is either measured or inferred according to the method proposed in Sec. II-D. Moreover, $t_h(f)$ is a measure indicating that the tendency of a client to handoff decreases with coverage and thus increases with frequency f:

$$t_h(f) = \frac{f - F_{min}}{F_{max} - F_{min}},$$

where F_{min}, F_{max} denote the minimum and maximum frequencies in the wide-spectrum network, respectively. Next, $H(k,f)$ denotes the delay incurred with the selection of AP k and carrier frequency f as a new association:

$$H(k,f) = \begin{cases} 0, & \text{if } k = AP_{cur}, \quad f = f_{cur} \\ \frac{D_{sw}}{D_{sw} + D_{AP}}, & \text{if } k = AP_{cur}, \quad f \neq f_{cur} \\ \frac{D_{AP}}{D_{sw} + D_{AP}}, & \text{if } k \neq AP_{cur}, \end{cases}$$

where AP_{cur} and f_{cur} denote the AP and frequency of the current association respectively, while D_{sw}, D_{AP} denote values for the channel switching and handoff delay, respectively. Finally, $d_{tr}(P)$ is a measure expressing the transmission delay under channel quality P:

$$d_{tr}(P) = 1 - \frac{R_i(P)}{\max_P\{R_i(P)\}},$$

where i is the band including carrier frequency f, and $\hat{P}_{k,f}(g_i)$ denotes the inferred or measured channel quality of the client's link to AP k at location g_i and frequency f.