IBM Research Report

Adapting Cellular Networks to Whitespaces Spectrum

Mukundan Madhavan
IBM Research India
Bangalore, India.

Harish Ganapathy
The University of Texas,
Austin, United States of America.

Malolan Chetlur
IBM Research India
Bangalore, India.

Shivkumar Kalyanaraman
IBM Research India
Bangalore, India.

IBM Research Division
Almaden - Austin - Beijing - Delhi - Haifa - T.J. Watson - Tokyo - Zurich

***Abstract*—TV whitespaces, recently opened up by the FCC for unlicensed use, are seen as a potential cellular offload and/or standalone mechanism, especially in dense metros where the demand for throughput is high. In this paper, we use real data collected from whitespaces databases to empirically demonstrate features unique to whitespaces - "power-spectrum trade-off" and spatial variation in spectrum availability. From this study, we conclude the need for whitespaces-specific adaptations to cellular networks so as to be able to extract maximum throughput and guarantee reliability. To tackle the effects of the power-spectrum trade-off, we propose a novel base-station design that specifically uses low-power transmitters as a means to maximise throughput. This design co-locates and networks together many low-powered mode-I devices to act as a multiple-antenna array. We estimate the size of the array required to meet typical rate targets, and show that the array design significantly outperforms traditional designs in terms of throughput for a given cost. We then turn our attention to spatial variability and study its impact on the problem of locating base-stations in a whitespaces network. Here, we propose spectrum-aware placement algorithms for whitespaces, which account for this spatial variability along with key parameters like user density. We show that such algorithms clearly outperform traditional placement algorithms and improve network coverage in this band.**

I. Introduction

The Federal Communications Commission (FCC) recently passed a ruling [2] that opens up *Television Whitespaces*, (the UHF band from 54 MHz to 806 MHz) for unlicensed use by wireless devices. The unlicensed devices are expected to function as "secondary devices" in these bands, i.e., they are expected to share the spectrum with primary occupants such as TV stations (and wireless microphones to a lesser degree) without causing interference to them. Later, the FCC called [3] for the presence of a central database with which primary devices can reserve spectrum for their exclusive use. These reservations detail the frequency bands they expect to occupy, their period of occupation, and their locations. The central database aggregates all primary reservations, creates and provides any querying unlicensed device with a *spectrum map* of TV bands (chunks of 6MHz spectrum) that the device can use at its location. The database is currently set to be updated once per day.

The FCC classifies unlicensed whitespace devices into three categories based on the power at which they can transmit: *(i) Fixed devices* can transmit at power levels upto 4W. They play a role in whitespace networks that is akin to cellular base-stations. *(ii) Mode-II devices* can transmit at powers upto 100mW. They play a role that is akin to smaller (WiFi) access points. *(iii) Mode-I devices* can transmit at powers upto 40mW. They are similar to mobile or other handheld devices in a cellular network.

A. Contributions and organization of paper

In this paper, we present real data collected from two whitespaces database pilots hosted by Spectrum Bridge [4] and Telcordia [5]. We derive new insights about whitespaces from this data and use these insights to propose novel design techniques for whitespace networks. Section II augments data from our work in [1], and presents the resultant observations. This section forms the backbone of our paper, and is a key contribution to whitespaces literature. Our observations in this section fall under two categories:

1) We observe the existence of a "power-spectrum trade-off", where low-powered mode-I devices have access to more spectrum than higher-powered fixed devices. We show that this difference in spectrum is often, quite significant. Importantly, we give insights on how this trade-off can be exploited to alleviate the severe lack of whitespaces spectrum in dense metros.
2) We demonstrate that the availability of whitespaces spectrum is a function of space, and changes over the scale of a few kilometers[1]. Further, the number of available TV bands can change multiple times with varying patterns over this scale of distance. Finally, in addition to such gradual large-scale variation, abrupt localized variations also exist in some dense metros.

As a result of these observations, we propose and analyze two modifications to a traditional network that help achieve higher throughputs and coverage. By traditional network, we mean a network with a single high-powered/expensive base-station serving many low-powered/inexpensive users/mobiles.

In Part I of the paper, we argue that large arrays[2] of *co-located* low-powered mode-I devices networked together are a *compelling, if not necessary*, base-station architecture in whitespaces because of the extra spectral diversity that comes into play by lowering power. The specific technical contributions in this part are:

3) in Section IV, we formulate the *optimal base-station composition* problem, i.e., the minimum number of low-powered mode-I devices required to achieve a given set of long-term throughput targets. We also demonstrate the usefulness of the max-sum-rate scheduler in long-term planning for networks.
4) We find the optimal array composition through numerical experiments under real-world settings in Section V. We demonstrate that the proposed design can support throughput-demanding applications such as high-quality movie streaming in a much more cost-efficient manner than a traditional design, which is often prohibitively costly in these regimes.

The scale of spatial variation observed in Section II leads us to believe that its impact is higher on network-wide problems rather than the design of a single base-station cell. Consequently, in Part II of the paper (Section VI onwards), we consider the problem of placing the large arrays designed in Part I over a given geographical area, as a function of the spatial variability of whitespaces spectrum. We argue that whereas traditional cellular node placement is primarily carried out based on the location of user hot spots, whitespaces node placement must additionally account for the spatial availability of spectrum. The specific technical contributions in the second part are as follows:

6) In Section VI, we present two different formulations of the whitespaces base-station placement problem.

[1]This is not the case in cellular networks where providers typically have access to the same spectrum over a distance of the order of a kilometer.

[2]Recently, there has been a growing belief that MIMO arrays comprised of a large number of cheap, small, antennas (called "Massive MIMO" by some) is the next step towards breaking the throughput barrier [6]–[9].

7) In Section VII, we prove that the objective functions in these formulations are *sub-modular*, a property that is useful in light of the fact that the proposed formulations are NP-hard. Low-complexity greedy algorithms are then presented that have an accuracy guarantee of $\left(1-\frac{1}{e}\right)$.
8) We then propose alternate solutions based on convex relaxations and the well-known k-clustering technique in Sections VIII and IX respectively. The convex relaxation technique is shown to be optimal under certain settings, and the clustering algorithm provides a useful alternative with a different complexity trade-off.
9) In Section X, we numerically evaluate these proposed algorithms and show that spectrum-aware placement can, in some cases, even double the average rate supported by the network compared to traditional placement. We also demonstrate how spectrum aware-placement is crucial in dense metros, where spectrum is relatively scarce.

In summary, we propose a shift from a "traditional" placement of high-power base-stations to low-power antenna-arrays that are placed in a spectrally-cognizant manner, as a means to achieve significant throughput gains in whitespaces.

B. Related work

In general, research concerning the design of database-driven whitespaces networks is still in its nascent stages. Murty et al. [10] architect a whitespaces network called *Senseless* that intelligently addresses the practical challenges proposed by a database/geo-location driven system. The authors propose and evaluate the use of a transmit beacon to enable the entry of a mode-I device, that has no prior knowledge of the spectrum map at its location, into the white space network. Further, they also propose an algorithm that adapts the resolution of the location information and study the associated performance tradeoffs.

In addition to engineering standalone whitespaces systems, other literature in the area is centered around the evolving IEEE 802.22 Regional Access Network standard [11], [12]. In particular, the topic of co-existence [13] amongst many 802.22 networks has received attention. Both cooperative [14], [15] and non-cooperative or game theoretic [16], [17] spectrum sharing protocols have been proposed and evaluated. An important take-away from the extensive empirical evidence provided in the above papers [10] and others [18] is that whitespaces spectrum maps can exhibit spatio-temporal variability.

II. Whitespaces - Unique Challenges and Solutions

A. Power-spectrum trade-off: A case for an array of mode-I devices

Consider the spectrum available at an FCC-registered cell tower [19], [20] in Middlesex County, New Jersey, USA, a densely-populated region twenty four miles west of New York City. Suppose that this tower operator is interested in equipping the tower with whitespaces capability. Figure 1(a) shows the number of TV bands available to a fixed device that would hypothetically replace or supplement the cell tower at this location, as seen on the whitespaces database [4]. As we can see, this fixed device has access to only one TV band (i.e., 6 MHz of bandwidth). In comparison, a standard WiMAX base-station operates using 10 MHz or 20 MHz of bandwidth per sector.

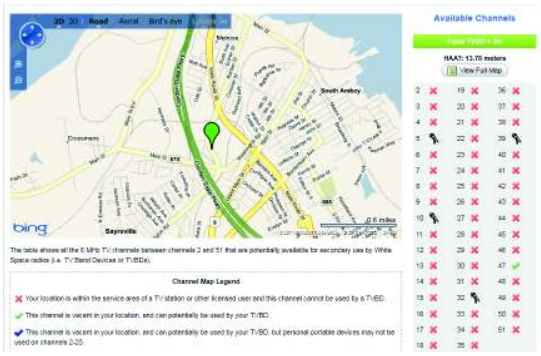

(a) Fixed device

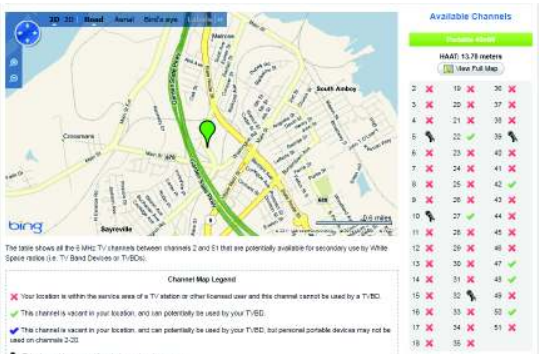

(b) Mode-I device

Fig. 1. Spectrum map [4] of devices at 2 Gowin St., Sayreville, NJ 08872.

Now, consider a mode-I device (to function as a base-station) placed at the *same* location. Figure 1(b) shows that this mode-I device has access to five TV bands (i.e., 30 MHz) instead of one. Thus, we see that lower-power mode-I devices can have significantly higher spectrum available to them than fixed devices at the same location and time. One possible reason for this is that they are less likely to cause interference to incumbent primary devices. Empirical evidence from [4] strongly suggests that this phenomenon of power-spectrum trade-off is widely prevalent, and not restricted to any one geography.

This effect has significant implications for network design, especially in dense cities. Dense cities tend to have more primary devices (TV transmitters), and consequently, lower availability of secondary-access spectrum. Consequently, the spectrum afforded to high-power fixed devices is low or even, non-existent. Consider the example of the Empire State Building in New York shown in Figure 2. Here, we see that the fixed device has access to *no spectrum at all*. However, a mode-I device placed at the same location has access to a bandwidth of 6 MHz, which allows for meaningful service.

Strangely, mode-II devices, which lie between fixed and mode-I devices in terms of power, do not find a place on the power-spectrum tradeoff curve. It is observed in the context of mode-II devices that the reduction in transmit power is just

(a) Fixed device

(b) Mode-I device

Fig. 2. Spectrum map [4] for devices at the Empire State building

not sufficient to release more spectrum. Figure 3 is a typical sample of our observations where mode-II devices have at most as many TV bands as a fixed device and hence offer no added spectrum.

Fig. 3. Spectrum map [4] for a mode-II device at 2 Gowin St., Sayreville, NJ 08872 [20].

Design implications: The use of mode-I devices for transmission would trade-off power for extra bandwidth. Considering that power is only a logarithmic component in capacity, whereas bandwidth is a linear factor, this seems like a beneficial choice. However, this design choice also immediately raises concerns about the hundred-fold loss in transmit power compared to a fixed device. In the high-SNR regime, it seems that this hundred-fold loss would be less debilitating, given the nature of the log function. However, at low-to-moderate SNRs, the additional spectrum is not always sufficient to overcome the power loss. Under settings where the additional spectrum is insufficient, we propose the use of an array of mode-I devices networked together to perform *transmit diversity*. We show that by exploiting spatial diversity in addition to the inherent extra spectrum, the array design significantly outperforms the traditional design that uses fixed devices in terms of cost required to meet the same throughput.

For example, our experiments (detailed in Section V) show that typically, a single mode-I device does not outperform a single fixed device (in terms of rate) except under very low interference conditions [3]. However, we find that, depending on the setting, the mode-I design consistently outperforms the fixed design when the size of the mode-I array is 5 to 20 times the number of fixed devices employed. Making a reasonable assumption that the cost of portable devices is akin to that of a mobile phone or router whereas the cost of a base-station is at least two orders larger [21], the array design has the potential to be more cost-effective. We also note that at higher rate settings (like 2 Mbps), the number of fixed devices required to meet the target are prohibitively large, thus making the fixed design infeasible. In such settings, the mode-I design is even more compelling. We do recognize however that there are costs and challenges associated with networking together co-located devices (e.g., synchronization) that are arguably more difficult to quantify and are beyond the scope of this paper.

We provide more formal analysis of this design in Sections IV and V, where we show how it comprehensively outperforms standard designs.

B. Heterogeneity of spectrum: The case for spatially-aware placement

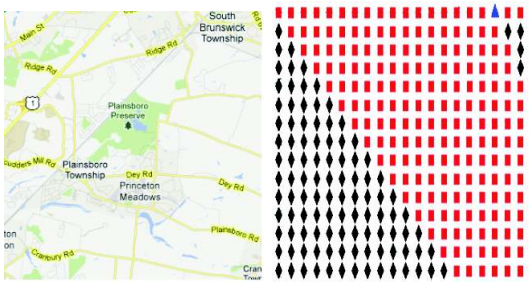

(a) 8 sq. km in New Jersey (b) Number of TV bands

Fig. 4. Number of TV bands available to mode-I devices in a region of New Jersey [5]. Color map : Black diamonds - 6 TV bands, Red squares - 7 TV bands, Blue triangles - 8 TV bands

Figure 4 shows the whitespaces spectrum available over a 8 sq. km area of New Jersey, USA, as shown by the Telcordia whitespaces database [5]. Over this area, we queried the database for the spectrum available to a mode-I device at intervals of 500 meters. The number of channels available at these various points are depicted in Figure 4(b), with different colors used to denote a different number of channels. Similar data from two other regions in Newark, USA and around the Rockefeller Center in New York are also shown in Figures 5 and 6 respectively. We make the following key observations from these figures:

[3]Both can however support low-rate requirements like VoIP for normal cell sizes of $1-2$ kms.

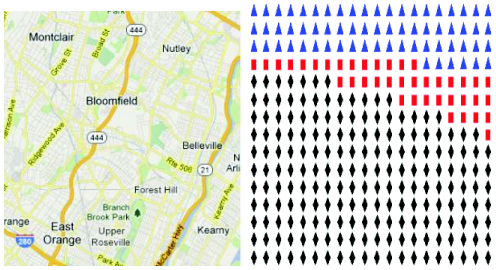

(a) 8 sq. km in Newark City (b) Number of TV bands

Fig. 5. Number of TV bands available to mode-I devices in a region of Newark [5]. Color map : Black diamonds- 1 TV band, Red squares - 2 TV bands, Blue triangles- 3 TV bands

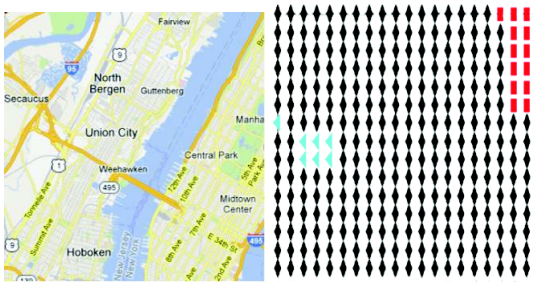

(a) 8 sq. km around Rockefeller Center (b) Number of TV bands

Fig. 6. Number of TV bands available to mode-I devices around Rockefeller Center [5]. Color map : Cyan triangles - 0 TV bands, Black diamonds - 1 TV band, Red squares - 2 TV bands

1) Spatial variation manifests itself as a gradual phenomenon over the scale of a few kilometers. At this scale, it is not always relevant to the design of a single base-station cell, except to those cells that fall in the transition regions. However, there are significant changes in spectrum within the scale of coverage of a single network, making it relevant to network-level design problems.
2) Over the scale of a network's coverage, the number of bands available can change multiple times. For example, in Figure 5, the number of bands available changes from three to two, and again transitions from two to one within a 5 *km* distance.
3) In addition to gradual changes over kilometers, abrupt and localized changes are also possible. This can be seen in the cyan points of Figure 6(b), where the spectrum changes abruptly from one to zero bands over a small region.
4) The change in spectrum does not follow any specific pattern across the three regions, making it difficult to model the phenomenon in a generic way. Consequently, spatial variation has to be factored in as a geography-specific input parameter into the design of networks.
5) Interestingly, we observed no temporal variations in spectrum over the ten-day window through which the data was collected. Though this is certainly not conclusive evidence about the lack of time variability in spectrum, it is indicative that presently, time-variability is either very gradual or not present in all regions.

Design Implications: As we can see, spatial variability occurs on a scale that is typically larger than a single cell, but smaller than a network's coverage area. While this variation can impact different aspects of network design, we feel that the problem of choosing base-station locations merits special attention. In traditional cellular networks, base-station placement typically considers parameters such as user hot spots and the desired coverage region. The following example illustrates why such a placement algorithm, which is agnostic to spatial variation in spectrum, will fail in whitespaces. Suppose one were to find a user hot spot in the cyan region of Figure 6(b). Placing a base-station at this location would be meaningless since it would have access to no spectrum. On the other hand, a base-station placed in the red region would be too far from this user hot spot to offer sufficient throughput. Thus, the better choice would be to balance these two requirements and place a base-station at an intermediate point. Later in the paper, we demonstrate through more detailed experiments, the importance of spectrum-aware placement, especially in dense cities, where whitespaces spectrum is relatively scarce.

In the past, there has been work on optimal placement problems in the related, but slightly different context of wireless sensor networks. These works consider placement methods that optimize various parameters like power-efficiency [22], coverage (e.g., [23], [24]). In the cellular context, Zhang and Andrews [25] study the performance of a random (Poisson) placement of distributed antennas under various diversity schemes and interference scenarios [25]. However, past literature has not dealt with variable spectrum availability, which is a key feature in whitespaces. In that sense, the traditional cellular placement problem is a special case of the placement problems formulated in this paper.

Lastly, we point out that while our analysis of placement problems in this paper does not model temporal variation, it can be readily extended to cover this case by using, for instance, time-averaged spatial distributions for spectrum. In further sections, we explore the two identified complementary design problems in greater detail. We first explain our system model briefly in the next section. This is followed by an analysis of our novel base-station design in Sections IV and V. Spectrum-aware placement of these novel base-stations is analysed beginning from Section VI.

III. System Model

We focus on the downlink of a network that operates in slotted-time. Since we are interested in the long-term design of networks (base-station design, placement etc), the network state variables often take the form of distributions and/or bounds that do not vary with time; examples include spatial distributions of users, the maximum number of interferers, cell capacity and traffic load targets, etc.

Modelling users, locations and traffic: Consider a whitespaces network with B base-station cells. Each base-station is assumed to have K' users that it must serve. Each user is represented using a mode-I device and has

an associated target rate requirement of λ' [4]. The base-station cell is assumed to be a square cell of side D *km*, and users are scattered over this square area according to some spatial distribution. The base-station is located at the centre $\left(\frac{D}{2}, \frac{D}{2}\right)$. For ease of analysis, we partition this cell into N microcells, where N is chosen to be a perfect square for simplicity. We index the micro-cells by tuples $\{(i,j)\}_{i,j=1}^{\sqrt{N}}$ beginning with the bottom-left corner. User locations are quantized to the center of each microcell given by $\frac{D}{\sqrt{N}}\left[(i-1)\ (j-1)\right]^T + \frac{D}{\sqrt{2N}}$.

Base-station composition: The base-station is composed of n_I mode-I devices networked together. Ideally, n_I should be chosen as the smallest value required to meet certain target throughputs in the system. The minimum-n_I required to meet a given desired rate-target λ' in a cell with K' users, and its comparison with traditional designs will be a subject of our study shortly.

Spectrum map: Let the set[5] of all whitespaces TV bands be denoted by $\mathcal{T} = \{2, 3, \ldots, 51\}$. The set of available TV bands at location l_i is given by $\mathcal{T}(l_i) \subseteq \mathcal{T} \setminus \{2, 3, \ldots, 21\}$. Supported by our empirical data, we also make the assumption that the spectrum availability is homogeneous within each cell. We also define $\mathcal{T}_F$ and $\mathcal{T}_I$ to denote the spectrum available to a fixed-device and a mode-I device respectively in the considered cell.

Multiple-access scheme: We assume that the base-station operates using orthogonal-frequency-division-multiple-access (OFDMA) over the set of TV bands $\mathcal{T}_I$ available to it. OFDMA is the multiple-access scheme of choice in standards such as IEEE 802.22 [11]. Each TV band of bandwidth $B = 6$MHz is sub-divided into L OFDMA sub-bands for transmission.

Transmit power: According to the FCC mandates, a fixed and mode-I device can transmit at a maximum power of $P_F = 4$W and $P_I = 40$mW respectively. For simplicity, we assume that on each sub-band f on TV band $b \in \mathcal{T}_I$, a base-station device type-d transmits at power $\frac{P_d}{L},\ d \in \{F, I\}$. Equal power transmissions are suitable for the whitespaces setting where a peak power constraint is imposed. On the other hand, rate-optimal approaches such as *waterfilling* are based on an average power constraint that might cause unacceptable levels of interference to the incumbents at any given instant.

Spatial distribution: Within the base-station cell, each user is thrown independently according to a distribution $\rho_i(l)$ with support $\left\{\left((i-1)\frac{D}{\sqrt{N}} + \frac{D}{\sqrt{2N}}, (j-1)\frac{D}{\sqrt{N}} + \frac{D}{\sqrt{2N}}\right)\right\}_{i,j=1}^{\sqrt{N}}$. The users' positions induce a path-loss gain [26] that is given by $\alpha_i = 10^{-(32.45+20\log_{10}(f_c)+20\log_{10}(d_i))}$, where d_i is the distance from the i-th user's location to the base-station at $[\frac{D}{2}\ \frac{D}{2}]$, f_c denotes the carrier frequency. While we have considered a free-space path-loss model here, we have also verified that the intuition derived holds for more complex models like the Sub-urban HATA model [27].

Small-scale Rayleigh fading and transmit diversity: Let h_{kfbn} be the discrete random variable used to denote the small-scale fading channel gain to user k on OFDMA sub-band f and TV band b on device copy $n = 1, 2, \ldots, n_I$ of the mode-I array; $h_{kfbn} \sim \exp(1),\ \forall k, f, b, n$, is exponentially-distributed. We assume that all devices of the base-station array co-operate by transmitting the same data, a technique popularly called *transmit diversity*. In view of not violating the peak power constraint imposed by the FCC mandates, we do not propose matching the antennas/devices to the channel, which clearly promises higher signal quality. In particular, these antennas are separated by at least half the wavelength, each transmitting the same data at power $\frac{P_d}{L}$.

This technique motivates the definition of an *effective* channel gain to user k on OFDMA sub-band f in TV band b as $h_{kfb}(\mathcal{N}) = \sum_{n\in\mathcal{N}} h_{kfbn},\ \mathcal{N} = \{1, 2, \ldots, n_I\}$. From standard probability theory, we know that $h_{kfb}(\mathcal{N}) \sim \chi^2(|\mathcal{N}|)$.

Interference model: We assume that there are $N_{F,int}$ and $N_{I,int}$ fixed and mode-I interferers respectively in the cell that are drawn from a uniform distribution $\rho_{int}(l)$ on the cell. The fixed and mode-I interferers are stationary and the path-loss gain between the j-th fixed (resp. mode-I) interferer and the k-th user is denoted by β_{kj}^F (resp. β_{kj}^p). The Rayleigh fading gain between the j-th fixed (resp. mode-I) interferer and the k-th user in sub-band f of TV band b is denoted by g_{kjfb}^F (resp. g_{kjfb}^p).

Scheduling and average rate: Let Ω denote the set of all possible schedulers. Then, for a base-station with n_I mode-I devices following a scheduling policy $s \in \Omega$, we denote the long-term average service rate by $\vec{\nu}^{ave}(s, n_I)$. The k^{th} element of this vector represents the long-term average rate of the k^{th} user, and is denoted by $\nu_k^{ave}(s, n_I)$. Searching over all scheduling policies $s \in \Omega$ characterizes the long-term rate region $\mathcal{V}_{LT}(n_I)$ of the system.

IV. Part I: Building the array base-station

Having proposed a base-station design that comprises of an array of mode-I devices, we now formulate the base-station composition problem where we determine the optimal composition (minimum-n_I) as a function of the long-term system parameters. This minimum-n_I is given by:

$$n_I^* = \begin{array}{ll} \arg\min & n_I \\ \text{s.t.} & \nu_k^{ave}(s, n_I) \geq \lambda',\ \forall k = 1, 2, \ldots, K'. \\ & n_I \in \{0, 1, 2, \ldots, \infty\}, s \in \Omega. \end{array} \quad (1)$$

In other words, we identify the smallest array-size such that all K' users receive above a threshold throughput target λ'. An important property of the rate region that implies the feasibility of the above problem is the following: Let $\vec{\nu}_{\max}(n_I)$ denote the element-wise maximum across all symmetric rate vectors

[4]In the absence of more specific information on per-user patterns in the long-term, we assume symmetric rate targets for all users. Further, given a set of asymmetric throughput targets, one can always compute the maximum of the targets and build the array to target this new symmetric set of targets.

[5]For simplicity in notation, we ignore the six bands that are reserved for wireless microphones.

contained in the rate region. Then, $\vec{\nu}_{\max}(n_I)$ is *strictly increasing* in n_I due to transmit diversity, albeit while exhibiting diminishing returns. Thus, there always exists a scheduler s^* and composition n_I^* that meets the throughput demand λ' for each user.

There are two main challenges to actually finding a solution to (1). Firstly, it would involve searching over the space of schedulers Ω, which seems infeasible at first sight. Secondly, given any composition n_I, we need to calculate the long-term service rate $\vec{\nu}^{ave}(s, n_I)$. To address both these issues, we use to our advantage, the fact that we are interested in a long-term design, and make the assumption that all users have *symmetric channel distributions*, i.e., $\rho_1(l) = \rho_2(l) = \ldots = \rho_{K'}(l) = \rho(l),\ \forall l$ for all users. This is reasonable because in the long-term, one might not possess more specific information pertaining to the position of each user in the system.

Under this setting, we consider the use of the well-known *max-sum-rate* scheduler. This scheduler selects the user with the strongest channel in each OFDMA sub-band and independent decisions are made on each sub-band. In the context of our long-term design goals, we argue that it is sufficient to base our design on the max-sum-rate scheduler in particular without having to search over the searching over the space of schedulers Ω. This is because the max-sum-rate scheduler is optimal for symmetric rate points and channels, a notion that we formalize in the proposition below. We provide only a sketch of the proof in the interest of space.

Proposition 1. *If any scheduler $s \in \Omega$ can achieve (dominate) a symmetric rate target of λ_{sym}, then so can the max-sum-rate scheduler.*

Proof: Refer Appendix. ∎

Next, we show that the max-sum-rate scheduler, herein denoted by $s^* \in \Omega$, addresses the second challenge of being able to evaluate the long-term service rate. We start by defining the signal-to-interference-plus-noise ratio (SINR) at user k on frequency f and TV band b as

$$\text{SINR}_{kfb}(\mathcal{N}) = \begin{cases} \frac{(P_I/L)\alpha_k h_{kfb}(\mathcal{N})}{I^p_{kfb}+I^F_{kfb}+N_o}, & b \in \mathcal{T}_I \cap \mathcal{T}_F \\ \frac{(P_I/L)\alpha_k h_{kfb}(\mathcal{N})}{I^p_{kfb}+N_o}, & b \in \mathcal{T}_I \setminus \mathcal{T}_F \end{cases} \tag{2}$$

where $I^F_{kfb} = \frac{P_F}{L}\sum_{j=1}^{N_{F,int}} \beta^F_{kj} g^F_{kjfb}$ and $I^p_{fkb} = \frac{P_F}{L}\sum_{j=1}^{N_{p,int}} \beta^p_{kj} g^p_{kjfb}$ are the portable and fixed interference powers and N_o is the additive noise power. In the interest of computational complexity, we consider a modified version of the the above SINR expression, as shown in (3), that contains the expected interference contributions instead of realizations in the denominator , i.e.,

$$\hat{\text{SINR}}_{kfb}(\mathcal{N}) = \begin{cases} \frac{(P_I/L)\alpha_k h_{kfb}(\mathcal{N})}{\bar{I}^p_{kfb}+\bar{I}^F_{kfb}+N_o}, & b \in \mathcal{T}_I \cap \mathcal{T}_F \\ \frac{(P_I/L)\alpha_k h_{kfb}(\mathcal{N})}{\bar{I}^p_{kfb}+N_o}, & b \in \mathcal{T}_I \setminus \mathcal{T}_F \end{cases} \tag{3}$$

where $\bar{I}^F_{kfb} = \mathbb{E}[I^F_{kfb}]$ and $\bar{I}^p_{kfb} = \mathbb{E}[I^p_{kfb}]$. Here, the expectations are computed over both small-scale fading channels and positions of the interferers. Note that the resulting variables $\bar{I}^F_{kfb}$ and $\bar{I}^p_{kfb}$ are now deterministic functions of the position of user k and hence random variables themselves.

The max-sum-rate scheduler schedules the user with the highest SINR on each OFDMA band and therefore the system-wide rate is given as

$$\begin{aligned} & \textstyle\sum_{k=1}^{K'} \nu^{ave}_k(s^*, n_I) \\ = \;& L\textstyle\sum_{b\in\mathcal{T}} \mathbb{E}\left[\log_2\left(1+\max_k \hat{\text{SINR}}_{kfbI}(\mathcal{N})\right)\right] \\ = \;& L\left(\textstyle\sum_{b\in\mathcal{T}_F\cap\mathcal{T}_I} \mathbb{E}\left[\log_2\left(1+\max_k \frac{(P_I/L)\alpha_k h_{kfb}(\mathcal{N})}{\bar{I}^p_{kfb}+\bar{I}^F_{kfb}+N_o}\right)\right]\right. \\ & \left. + \textstyle\sum_{b\in\mathcal{T}_I\setminus\mathcal{T}_F} \mathbb{E}\left[\log_2\left(1+\max_k \frac{(P_I/L)\alpha_k h_{kfb}(\mathcal{N})}{\bar{I}^p_{kfb}+N_o}\right)\right]\right) \end{aligned} \tag{4}$$

where we have used the fact that the SINRs on all OFDMA sub-bands are statistically equivalent. The expectation here is over the small-scale fading channels and positions of the in-network users. The per-user rate follows easily from (4) since each user receives the same rate in a symmetric setting. Thus

$$\nu^{ave}_k(s^*, n_I) = \frac{\sum_k \nu^{ave}_k(s^*, n_I)}{K'}. \tag{5}$$

The above expression can be estimated efficiently through Monte Carlo methods. We re-iterate that the only randomness remaining in the system is that induced by the small-scale and large-scale fading channels. When conditioned on the positions of the users, the term $\hat{\text{SINR}}_{kfb}(\mathcal{N})$ is essentially a maximum over K' non-identical chi-squared random variables of order n_I. By substituting (5) into (1), we now have a form that is optimizable. In the next section, we solve the optimal composition problem under many realistic network settings.

V. Optimal Composition

In this section, we evaluate the approach presented in the last section using network scenarios that are inspired by realistic applications. To illustrate how prohibitive the fixed design is in these settings, we also estimate the size of an array of high-powered fixed devices needed to meet these rates for all the users, when a single fixed device is insufficient.

A. Simulation setup

We vary the following long-term state parameters and solve (1) in each case: (K', λ'), the spatial distributions,$\{\rho_i(l)\}$, the cell size D and the number of interferers $N_{F,int}$ and $N_{p,int}$. The simulation models the cell tower in New Jersey (Figure 1(b)) with $\mathcal{T}_F = 47$ and $\mathcal{T}_I = \{22, 27, 42, 47, 48, 50\}$. We focus on a movie streaming application as defined in [28] where the targeted cell capacity and application bit-rate are given by $K' = 10$ users, $\lambda' = 2$ Mbps.

We now describe in detail how we choose the the remaining long-term state parameters in our experiments:

1) *Spatial distributions*: We consider two types of user distributions on the square cell. First, we consider the standard uniform distribution given by $\rho_i(l) = \frac{1}{N}, \forall i, l$, which essentially models the setting where we have no extra information on user positions. Next, we study the optimal composition under a cell-edge distribution that is given by $\rho_i(l) = \frac{1}{4\sqrt{N}-4}$ when $i, j = 1$ or $i, j = \sqrt{N}$ and $\rho_i(l) = 0$ otherwise. Such an edge-intensive distribution would represent the worst-case for design approach.

2) *Interference modelling and cell sizes*: We model fixed interference by considering cellular towers that are operating in the same area – potentially by different operators – over the whitespace bands (in standalone and/or offload mode). In particular, we are interested in WiMAX-order cell sizes $D \in \{1 \text{ km}, 2 \text{ km}, 5 \text{ km}\}$ and find that there are three, four and six towers/networks within these respective cell sizes (Figure 7(b)). Assuming that approximately one half of these towers are in uplink mode, we set $N_{F,int} = 1$, $N_{F,int} = 2$ and $N_{F,int} = 3$ for $D = 1$ km, $D = 2$ km and $D = 5$ km respectively. In the absence of more specific information, we treat the number of portable interferers per fixed interferer as a parameter. We analyze the optimal composition under different values of this ratio $\frac{N_{I,int}}{N_{F,int}} \in \{10, 20, \ldots\}$. For e.g., when $D = 5$ km with $N_{F,int} = 3$ fixed interfering towers, $\frac{N_{I,int}}{N_{F,int}} = 10$ means that we have $\frac{N_{I,int}}{N_{F,int}} \times N_{F,int} = 30$ portable interferers.

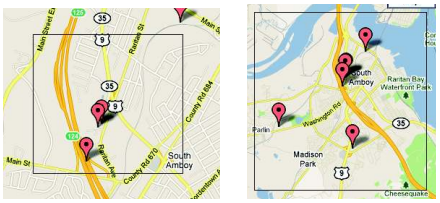

(a) 1 km cell (b) 5 km cell

Fig. 7. Number of fixed interferers [20] around 2 Gowin St., Sayreville, NJ 08872. Red markers denote cell-towers.

We compute minimum-n_I for this rate target, for the following network configurations $\{1 \text{ km}, 2 \text{ km}, 5 \text{ km}\} \times \{\text{uniform, cell-edge}\}$, where the per-user throughputs are averaged over 10000, 10000 and 5000 realizations of user positions, Rayleigh fading and interferer positions respectively.

B. Results on array sizes

In Figure 8(a), we plot the optimal composition (minimum number of mode-I devices) for the movie streaming case over the different network settings $(D, \rho_i(l)) \in \{1 \text{ km}, 2 \text{ km}, 5 \text{ km}\} \times \{\text{uniform, cell-edge}\}$ under the low-interference regime of $\frac{N_{I,int}}{N_{F,int}} = 10$. For comparison purposes, in the same figure, we also plot the minimum size of a competing design with only fixed devices constituting the array. The results show that for a cell radius of $D = 1$ km, under the cell-edge distribution, we need either 175 mode-I devices or 35 fixed devices to meet the specified rate target ($5\times$). Similarly, for a cell radius of $D = 5$ km under the cell-edge distribution, we need either 325 mode-I devices or 150 fixed devices ($2.2\times$). Operating under the premise that mode-I and fixed devices are akin (cost-wise) to household routers and base-stations respectively, the use of a large array of portable devices clearly leads to significant cost savings. Arguably, the fixed device design is not practical in these scenarios.

Admittedly, offered traffic models are not always reliable metrics because they are subject to change with space and time. Consequently, to give a more comprehensive comparison of the two designs, we have plotted the rate distributions resulting from each. We note that for the $D = 1$ km cell having

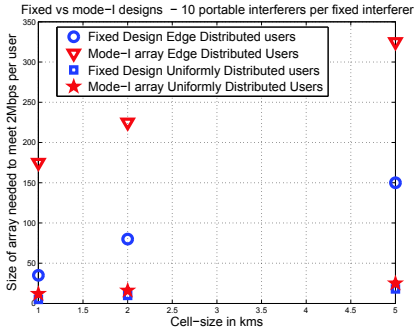

(a)

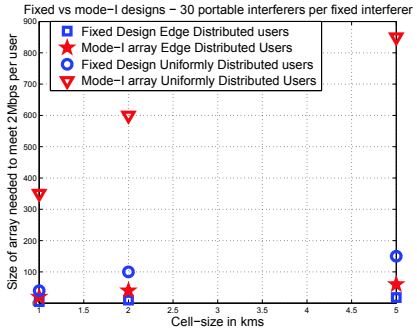

(b)

Fig. 8. Number of fixed and portable devices required to meet rate target of 2Mbps under (a) low interference $\frac{N_{I,int}}{N_{F,int}} = 10$, (b) moderate interference $\frac{N_{I,int}}{N_{F,int}} = 40$

uniform distribution of users, the 2 Mbps per user target is achieved by either 4 fixed devices or 12 mode-I devices with $\frac{N_{I,int}}{N_{F,int}}$ picked to be 10 (see Figure 8(a)).

Figure 9 below shows the cumulative distributions of rate across the cell for both these configurations. We can see that for both cases, the top 5 percentile users have identical rates (the plots are similar beyond the 95^{th} percentile). However, the fixed array has a better distribution up to this point. Since the uniform distribution is less likely to put many users near the cell edge, the 12 mode-I device array still performs as well as the 4-device fixed array.

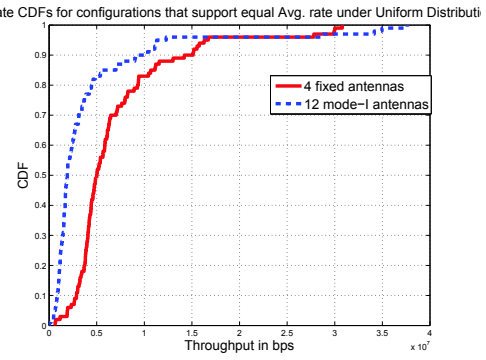

Fig. 9. Cumulative distribution of the rates across a cell of size $D = 1km$, $\frac{N_{I,int}}{N_{F,int}} = 10$, fixed array of size 4, portable array of size 12.

However, as we increase the mode-I array size beyond 12, the mode-I arrays' rate CDF quickly catches up to that of

the fixed array. For example, Figure 10 shows the Rate CDF comparison of 4 fixed devices against mode-I arrays of size 30 and 35 . We see that the CDF of the 35 device array is consistently better than the 4-device fixed array. (The optimal composition for the edge-distribution for these parameters is 175 mode-I devices.)

So, we summarise that for a typical real-life scenario, where a base-station is located taking user densities into consideration, a 12 device array is likely to be sufficient. For the mode-I array to do better consistently across the cell, a number closer to 35 will still suffice.

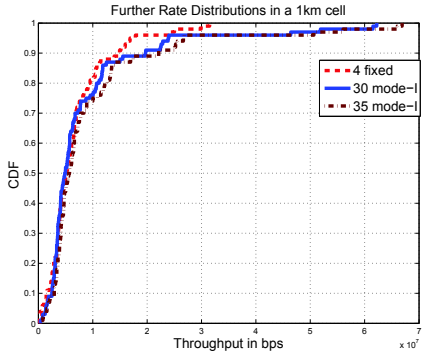

Fig. 10. Cumulative distribution of the rates across a cell of size $D = 1km$, $\frac{N_{I,int}}{N_{F,int}} = 10$, fixed array of size 4, portable array of size 30 and 35.

We continue this analysis in Figure 8(b) for the moderate interference regime of $\frac{N_{I,int}}{N_{F,int}} = 40$. Here, we see that under the cell-edge distribution, we need $10\times$ the number of portable devices for $D = 1$ km and $6.5\times$ for $D = 5$ km. To capture this trend, in Figure 11, we show the ratio of portable to fixed array sizes as a function of $\frac{N_{I,int}}{N_{F,int}}$, i.e., for various interference regimes for the $D = 1$ km case. Here, we see that under the cell-edge distribution and in the high interference setting of $\frac{N_{I,int}}{N_{F,int}} = 100$, we still only need at most 22 times the number of portable devices to meet the specified rate target.

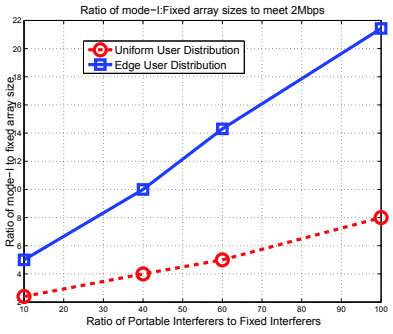

Fig. 11. Ratio of portable to fixed array sizes as a function of $\frac{N_{I,int}}{N_{F,int}}$ for $D = 1$ km.

Figure 12 shows a similar ratio of mode-I array size to fixed array size for the case of the sub-urban HATA [27] model. As expected, using this model increases the number of devices required to meet 2 Mbps rate under both designs. However, we find that the ratio of mode-I to fixed devices actually comes down under this model. This is because, as the array size required under the fixed design increases, the spectral diversity gain becomes saturated because the fixed design has only one TV band to operate on. Figure 3 below shows the ratio of mode-I array size to fixed array size for the HATA path-loss model for a 1 km cell under uniform spatial distribution. When we compare these results to the ones in Figure 9 of the revised paper, we see that as scattering comes into play, the extra spectrum from mode-I arrays is even more valuable. As a result, the mode-I design becomes even more cost-efficient in this model when compared to the fixed design.

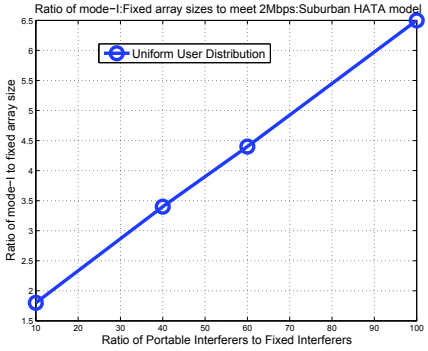

Fig. 12. Ratio of mode-I array size to fixed array size for $D = 1$ km, under Sub-urban HATA model.

.

Thus far, we have considered the effects of transmission power on whitespaces spectrum, and studied a novel base-station design to tackle this phenomenon. Having demonstrated the merits of this array design, we now move on to the second part of the paper, where we study the problem of base-station placement under heterogeneous spectrum.

VI. Part II- Base-station placement in space

In this part, we consider placing the mode-I array base-stations designed in Part I optimally over a given geographical area. We formulate the placement problem for whitespaces, provide solutions with varying complexities and accuracy, and evaluate these solutions using the collected whitespaces data. First, we make some additional notes on system modeling specific to the placement problem, and also make some minimal modifications to our previous notation to keep it consistent with the change of scale from a single cell to an entire network.

A. Some notes on modelling and notation:

Base-stations and users: In accordance with our design from Part-I, we assume that each base-station is made of an array of mode-I devices of size M. To maintain notational simplicity, we keep M fixed across all the base-stations. Users are also assumed to be mode-I devices as before. The spatial distribution of users in the i^{th} micro-cell of the coverage region is denoted using ρ_i.

Node placement: We assume that the network contains B base-stations. Earlier, we denoted the length of a base-station cell using D. However, D now represents the side of the *entire* network coverage region. As before, this network coverage region, is tessellated into N micro-cells for tractability. All base-station locations are quantised to the centers of

these micro-cells, indexed by the set $\theta = \{1, 2, .., N\}$. We note that the network coverage region, being larger than a cell, is divided into a larger number of micro-cells than before.

Channel modeling: We alter the notation used for path-loss from before. In Section III, we had used α_i to denote the path-loss from micro-cell i to a base-station located at the center of the square cell. However, in the placement problem, we must deal with multiple base-stations located within a square coverage area. Consequently, we use two subscripts for the pathloss and use α_{ij} to denote the pathloss from a user i to a base-station j.

The small-scale fading coefficient experienced by user i when receiving from antenna element y of base-station j on whitespaces band n is represented by the random variable g_{injy}. Similar to the pathloss co-efficient, the index j has been introduced to account for the existence of multiple base-stations in this problem. In accordance with the principles of the array design, these antennas exploit transmit diversity in their transmissions. Hence, the effective channel from the antenna array base-station to the user is denoted by h_{inj} where $|h_{inj}|^2 = \Sigma_{y=1}^{M} \mid g_{injy} \mid^2$, $\mid h_{inj} \mid^2 \sim \chi^2(M)$.

For each micro-cell (user) i, we also define $S_i = \frac{P_I/L}{I_i + N_O}$, where I_i denotes the cumulative interference from fixed and portable devices at location i. α_{ij}, h_{inj} and S_i together completely define the channel model in this setting.

Spectrum map: We recall that $\mathcal{T}$ denotes the set of all TV bands and that $\mathcal{T}(i)$ denotes the bands available at location i. We additionally define $\mathcal{A}_n(\mathcal{P}) = \{p \in \mathcal{P} : n \in \mathcal{T}(p)\}$, for $n \in \mathcal{T}$. $\mathcal{A}_n(\mathcal{P})$ denotes the subset of base-stations from base-station set $\mathcal{P}$ that have access to band n.

Diversity schemes: Given a placement of base-stations $\mathcal{P} \subseteq \theta$, we consider two types of popular diversity schemes. Under maximal ratio transmission (MRT), a user is served by all base-stations together, while under antenna selection (AS), a user is served by the base-station with the strongest channel [30]. Intuitively, the MRT scheme represents the best single-user rate that can be provided by a particular placement of stations, whereas the AS scheme is similar to a traditional network, where users are supported by the best out of the available cells. The achievable network-averaged spectral efficiency with MRT for a placement of antennas $\mathcal{P} \subseteq \theta$ may be expressed as

$$R_{MRT}(\mathcal{P}) = \sum_{i=1}^{N} \rho_i \sum_{n \in \mathcal{T}} \mathbb{E}\left[\log_2\left(1 + S_i \sum_{p \in \mathcal{A}_n(\mathcal{P})} \alpha_{ip}|h_{inp}|^2\right)\right]. \tag{6}$$

Note that we have exploited the fact that all sub-bands experience statistically equivalent fading to simplify our notation above. An equivalent expression that will be useful in further analysis is

$$R_{MRT}(\vec{\phi}) = \sum_{i} \rho_i \sum_{n \in \mathcal{T}} \mathbb{E}\left[\log_2\left(1 + S_i \sum_{p=1}^{N} \phi_p \Delta_{ipn} \alpha_{ip}|h_{inp}|^2\right)\right]. \tag{7}$$

Here, $\vec{\phi} = \{\phi_1, ..., \phi_N\}$ is a boolean vector, with $\phi_p = 1$ if a base-station is placed at location p and $\phi_p = 0$ otherwise. $\{\Delta_{ipn}\}_{i,p}$ represents a boolean matrix that encodes the spectrum map, i.e, $\Delta_{ipn} = 1$ if locations i and p share TV band n. Similarly, under AS diversity, the average rate is given by

$$R_{AS}(\vec{\phi}) = \sum_{i} \rho_i \sum_{n \in \mathcal{T}} \mathbb{E}\left[\log_2\left(1 + S_i \max_{p \in \theta} \phi_p \Delta_{ipn} \alpha_{ip}|h_{inp}|^2\right)\right]. \tag{8}$$

We observe that when the network-size is $B = 1$, both AS and MRT yield the same rate, i.e., $R_{AS}(\{p\}) = R_{MRT}(\{p\}),\ \forall p \in \theta$.

B. Optimal Placement Problem

We consider two formulations of the base-station placement problem:

- Given a budget of B antennas, finding a placement that maximizes the average network rate [6] in (6). This is hereafter referred to as the "max-rate" formulation.
- Determining the "smallest" placement set that guarantees a certain minimum target average network rate. This is hereafter referred to as the "min-cost" formulation.

These two formulations are standard in the wireless resource allocation literature (e.g., max-sum-rate vs. min-power subject to SNR constraints) and either may be used by a network designer as needed. Under the max-rate formulation, we are given a budget of B antennas and solve

$$\mathcal{P}^* = \arg \max_{\mathcal{P} \subseteq \theta,\ |\mathcal{P}| \leq B} R_{MRT/AS}(\mathcal{P}). \tag{9}$$

Here, $R_{MRT/AS}(\cdot)$ is used to convey that one may substitute either diversity scheme here. Under the min-cost formulation, we determine the minimum budget required to satisfy a given lower bound λ' on the average rate as follows:

$$\begin{aligned} \mathcal{P}^* \quad &= \quad \arg\min_{\mathcal{P} \subseteq \theta} |\mathcal{P}|. \\ &\text{s.t.} \quad R_{MRT/AS}(\mathcal{P}) \geq \lambda' \end{aligned} \tag{10}$$

We assume that λ' is such that (10) is feasible, i.e., $R_{MRT/AS}(\theta) \geq \lambda'$. While the algorithms and analyses proposed in this paper are applicable to both these formulations, we focus on the max-rate formulation in the interest of space. The closest analogue to our proposed formulations may be found in sensor placement literature, (see [31] and references therein) where they are called maximization and coverage problems respectively. However, the notion of utility in our context differs from that used in sensor placement [31], where the emphasis is on picking a subset of existing sensors to maximise the information obtained about all locations. In contrast, our problem is a more standard Shannon-capacity-based formulation.

It has been shown that sensor placement problems of this nature are typically NP-hard. While we do not formally prove hardness of our formulation in this paper, we observe that any brute-force approach to solving (9) and (10) would incur exponential complexity in the number of micro-cells N. In the next section, we appeal to the concept of sub-modular functions [35] to reduce the solution complexity while providing guarantees on its accuracy.

[6] In Sections VI - IX, we use the terms rate and spectral efficiency interchangeably.

VII. Leveraging Sub-modularity for tractable solutions

Towards designing tractable solutions to the placement problems, we first establish that the rates $R_{MRT}(\mathcal{P})$ and $R_{AS}(\mathcal{P})$ possess a special structure called *sub-modularity* that allows for efficient optimization. As a result, we show that (9) and (10) can be solved incurring polynomial complexities of $\mathcal{O}(KN^2W_{\max}|\mathcal{T}|)$ and $\mathcal{O}(N^3W_{\max}|\mathcal{T}|)$ respectively using the greedy algorithms proposed in Section VII-B, where $W_{\max}$ is the maximum number of TV bands available to any one user. We begin this section with a quick primer on sub-modular optimization (summarized from [33]–[35]).

A. Primer on sub-modularity

Let $\mathcal{E}$ be a finite set and $2^{\mathcal{E}}$ represent all its subsets. Then, $F: 2^{\mathcal{E}} \rightarrow \mathbb{R}_+$ is a *non-decreasing*, *normalized*, *sub-modular* function if *(i)* $F(\emptyset) = 0$ (normalized), *(ii)* $F(\mathcal{A}) \leq F(\mathcal{B})$ when $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{E}$ (non-decreasing) and *(iii)* $[F(\mathcal{A} \cup \{e\}) - F(\mathcal{A})] \geq [F(\mathcal{B} \cup \{e\}) - F(\mathcal{B})],\ e \in \mathcal{E} \setminus \mathcal{B},\ \forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{E}$ (sub-modular). Lemma 1 lists a relevant property of sub-modular functions.

Lemma 1. *If $F_n,\ n = 1, \ldots, N$, are sub-modular on set $\mathcal{E}$, then $\sum_{n=1}^{N} w_n F_n(\mathcal{A}),\ \mathcal{A} \subseteq \mathcal{E}$ is a sub-modular function for $w_n \geq 0, \forall n$.*

The optimization problems that have been often considered in the context of sub-modular functions are

$$F^* = \max_{\mathcal{A} \subseteq \mathcal{E},\ |\mathcal{A}| \leq K} F(\mathcal{A}) \tag{11}$$

and

$$F^* = \begin{array}{ll} \max_{\mathcal{A} \subseteq \mathcal{E}} & |\mathcal{A}| \\ \text{s.t. } F(\mathcal{A}) \geq c & \end{array}, \tag{12}$$

for some $0 < K < |\mathcal{E}|$ and $c > 0$. Since many NP-hard problems can be reduced to above forms, significant research has focused on developing efficient approximation algorithms for them. In particular, the performance of the greedy algorithm – that serially selects a location that maximizes the incremental cost has been extensively studied. Nemhauser et al. [36] considered problem (11) and showed that the greedy algorithm provides a $(1 - \frac{1}{e})$ approximation factor for this case. Some years later, Wolsey [32] considered problem (12) and showed that the greedy algorithm provides a $(1 + \log \max_{s \in \mathcal{E}} F(s))$ approximation factor for this case.

We refer the reader to Goundan et al. [33], Calinescu et al. [34] and Vondrak [35] for a summary of related results on sub-modular function optimization over other types of constraints. Next, we prove that $R(\mathcal{P})$ is a sub-modular function over $\mathcal{P} \subseteq \theta$.

B. Sub-modularity of $R(\mathcal{P})$ and the greedy algorithm

The following theorem constitutes one of the theoretical results of this paper.

Theorem 1. (i) *$R_{MRT}(\mathcal{P})$ is a sub-modular function over $\mathcal{P} \subseteq \theta$.*
(ii) *$R_{AS}(\mathcal{P})$ is a sub-modular function over $\mathcal{P} \subseteq \theta$.*

Proof: Refer Appendix. ∎

The seminal works by Nemhauser et al. [36] and Wolsey [32] have studied the performance of the greedy algorithm in solving subset selection problems such as (11) and (12), where the objective function is sub-modular. The greedy algorithms in the context of our node placement problem are given below.

Algorithm 1 Greedy node placement for max-rate (9)

1: Set $\mathcal{P} = \emptyset$.
2: **while** $|\mathcal{P}| \leq B$ **do**
3: Choose a location that maximizes

$$e^* = \arg \max_{e \in \theta \setminus \mathcal{P}} R_{MRT/AS}(\mathcal{P} \cup e) - R_{MRT/AS}(\mathcal{P}). \tag{13}$$

4: Set $\mathcal{P} = \mathcal{P} \cup e^*$.
5: **end while**
6: Set $\mathcal{P}_{greedy} = \mathcal{P}$.

Algorithm 2 Greedy node placement for min-cost (10)

1: Set $\mathcal{P} = \emptyset$.
2: **while** $R(\mathcal{P}) < R_t$ **do**
3: Choose a location that maximizes

$$e^* = \arg \max_{e \in \theta \setminus \mathcal{P}} R_{MRT/AS}(\mathcal{P} \cup e) - R_{MRT/AS}(\mathcal{P}). \tag{14}$$

4: Set $\mathcal{P} = \mathcal{P} \cup e^*$.
5: **end while**
6: Set $\mathcal{P}_{greedy} = \mathcal{P}$.

The following are celebrated results by Nemhauser et al. [36] and Wolsey [32]. We state them here without proof.

Theorem 2. *(i) In solving* (9)*, Algorithm 1 is $(1 - 1/e)$-optimal. (ii) In solving* (10)*, Algorithm 2 is $(1 + \log \max_{\mathcal{P} \in \theta: |\mathcal{P}| = 1} R(\mathcal{P}))$-optimal.*

Note that in the case of (10), the approximation factor depends on the objective function. In our case, $\max_{\mathcal{P} \in \theta: |\mathcal{P}| = 1} R(\mathcal{P})$ represents the maximum rate that may be obtained if we had a budget of only one antenna or node. As noted earlier, single node rates are the equal under AS and MRT. The following proposition quantifies this rate exactly.

Proposition 2. *The solution to the maximum rate single node placement problem under either MRT or AS is given by*

$$p^* = \arg \max_{p \in \theta} \sum_{i=1}^{N} \rho_i |\mathcal{T}(i) \cap \mathcal{T}(p)| \mathbb{E}\left[\log_2\left(1 + S_i \alpha_{ip} |h_{inp}|^2\right)\right]$$

Proof: Refer Appendix. ∎

Proposition 2 captures the trade-off that forms the main motivation of the second part of this paper- that maximum benefit may be achieved in whitespaces if base-stations are placed at positions that offer *a combination of high spectrum as well as proximity to user hot spots*. Finally, we study the computational complexity of the greedy approach in Proposition 3.

Proposition 3. (i) *Algorithm 1 has a worst-case complexity of $\mathcal{O}(BN^2W_{\max}|\mathcal{T}|)$ and* (ii) *Algorithm 2 has a worst-case complexity of $\mathcal{O}(N^3W_{\max}|\mathcal{T}|)$, where $W_{\max}$ is the maximum number of TV bands available to any one user.*

Proof: Refer Appendix. ∎

VIII. Approximation algorithms: Convex relaxation

In this section, we study the use of an alternate convex relaxation approach to solve (9) and (10). While this approach may, in general, be treated as an alternate approximation algorithm, we show that it is in fact, optimal in the low-SINR regime for MRT diversity. The standard approach in a convex relaxation is to replace the integer constraints in a combinatorial problem by some suitably chosen convex constraints. When the objective function is also convex or suitably convexified, this relaxed optimization problem may be solved efficiently. The final critical step is to round the fractional solution to obtain a feasible integer solution. By using the alternate expression for rate in (7), the equivalent integer problem to (9) can be written as

$$\vec{\phi}^* = \arg \max_{\vec{\phi}\in\{0,1\}^{|\theta|},\ \sum_l \phi_l \leq B} R_{MRT/AS}(\vec{\phi}). \tag{15}$$

For convenience, we introduce the vector $\vec{c}(i,n) = [\Delta_{i1n}\alpha_{i1}|h_{nl1}|^2\ \Delta_{i2n}\alpha_{i2}|h_{nl2}|^2 \dots \Delta_{iNn}\alpha_{iN}|h_{nlN}|^2]^T$, and concisely re-write the rate $R_{MRT}(\vec{\phi})$ as

$$R_{MRT}(\vec{\phi}) = \sum_i \rho_i \sum_{n\in\mathcal{T}} \mathbb{E}\left[\log_2\left(1 + S_i(\vec{\phi}^T\vec{c}(i,n))\right)\right]. \tag{16}$$

Note that $R_{AS}(\vec{\phi})$ takes the same form as $R_{MRT}(\vec{\phi})$ with the inner summation replaced by a maximum. We replace the boolean integer constraints in (15) with an interval and propose the following relaxation:

$$\vec{\phi}^*_{LP} = \arg \max_{\vec{\phi}\in[0,1]^{|\theta|},\ \sum_l \phi_l \leq B} R_{MRT/AS}(\vec{\phi}). \tag{17}$$

For the MRT case, $\log_2\left(1 + S_i\sum_{p=1}^N \phi_p\Delta_{ipn}\alpha_{ip}|h_{inp}|^2\right)$ is a logarithm of a linear function, which is a concave function, and concavity is preserved under a sum. Coupled with the linear constraints, (17) is clearly concave. However, the relaxation approach is not effective in the AS diversity case. The two natural relaxations that are possible in the AS setting can be obtained from the following two equivalent expressions of the objective function in (15). These are

$$\log_2\left(1 + S_i \max_{p\in\theta} \phi_p c_p(i,n)\right) = \max_p \log_2\left(1 + S_i\phi_p c_p(i,n)\right)$$

and

$$\log_2\left(1 + S_i \max_{p\in\theta} \phi_p c_p(i,n)\right) = \max_p \phi_p \log_2\left(1 + S_i c_p(i,n)\right).$$

The former re-expression yields a piecewise-concave relaxation while the latter yields a convex objective function. Unfortunately, neither relaxation is useful for efficient computation, and in fact, the latter relaxation can be shown to be NP-hard. In the next section, we propose an alternate heuristic based on clustering that is more suitable for the AS diversity problem.

Interestingly, in the case of the min-cost formulation in (10), the relaxation leads to a concave program under either diversity scheme as shown below.

$$\begin{array}{lll} \vec{\phi}^*_{LP} = & \arg\min_{\vec{\phi}\in[0,1]^{|\theta|}} & \sum_{p\in\theta}\phi_p \\ & \text{s.t.} & R_{MRT/AS}(\vec{\phi}) \geq R_t, \end{array} \tag{18}$$

(17) and (18) can be solved in polynomial-time using standard interior point methods [37]. Once the fractional solution $\vec{\phi}^*_{LP}$ is computed, the integer solution $\{\vec{\phi}_{IP}\}$ is generated by setting the largest B values of $\vec{\phi}^*_{LP}$ to one and the remaining values/locations to zero. While a universal guarantee on the accuracy of this approach is difficult to obtain, the following theorem characterizes the accuracy guarantee of (17) in solving (15) in the case of MRT diversity and under a *low-SINR regime*, where high interference prevails.

Theorem 3. *The convex relaxation in (17) solves (15) exactly in the case of MRT diversity under a low-SINR regime. Let $\{i_k\}_{k=1}^B$ index the maximum B values from the set $\left\{\sum_i \rho_i S_i \sum_{n\in\mathcal{T}} \mathbb{E}\left[\vec{c}(i,n)\right]\right\}_{i=1}^N$. Then, the optimal solution is given by*

$$\phi_j^* = \begin{cases} 1, & j \in \{i_k\}_{k=1}^B \\ 0, & \textit{else}, \end{cases} \tag{19}$$

for $j = 1, 2, \dots, N$.

Proof: Refer Appendix. ∎

Thus far, we have proposed two approaches to solve (9) and (10). The earlier greedy approach may be applied either diversity scheme. Then, we proposed an approximate solution based on convex-optimization, which optimally solves the problem at low SINRs, but is applicable only under MRT diversity. In the next section, we present an alternate clustering-based heuristic that solves the problem for the AS case. Later, in Section X, we experimentally evaluate all the proposed algorithms and compare them to a spectrum-agnostic placement approach.

IX. Lower-complexity algorithms: Clustering

As seen from Proposition 3, the complexity of the greedy approach becomes higher in spectrum-rich locations. In this section, we present an alternative approach, whose complexity is independent of spectrum availability. This algorithm provides the network designer with an alternative complexity trade-off to the greedy approach. The algorithm presented here finds its roots in the theory of expectation-maximization [38]. It has found application in many other areas such as vector quantization where it is popularly known as *Lloyd's algorithm* [39] and in data mining, where it is often called *k-clustering* [40].

We begin the algorithm with an arbitrary placement of B base-stations. We determine a "Voronoi" partitioning or clustering of the grid, i.e., for each location i, associate it with the base-station from which it has maximum rate. Next, given this Voronoi clustering, re-distribute or re-compute the "optimal" base-station location *within each cluster*. These steps are repeated until the increase in (9) diminishes below a threshold. The algorithm is stated formally below.

The following proposition argues that the above algorithm indeed terminates. Since theorems of this form are standard in clustering literature, we only provide a sketch of its proof in the interest of space.

Proposition 4. *Algorithm 3 terminates in a finite number $0 < T_\varepsilon < \infty$ of iterations.*

Proof sketch: Refer Appendix. ∎

Algorithm 3 Clustering-based placement for AS diversity

Pick an arbitrary initialization of positions $\vec{\phi}_{new} \in \{0,1\}^{|\theta|}$, $\sum_l \phi_l \leq B$, set $\vec{\phi}_{old} = [0\ 0 \ldots 0]^T$ and choose a suitable $\varepsilon > 0$.
while $(R_{AS}(\vec{\phi}_{new}) - R_{AS}(\vec{\phi}_{old}) \geq \varepsilon)$ **do**
  Set $\vec{\phi}_{old} = \vec{\phi}_{new}$ and reset $\vec{\phi}_{new} = [0\ 0 \ldots 0]^T$.
  Initialize clusters $\mathcal{C}_k = \emptyset,\ \forall k = 1, 2, \ldots, B$.
  for $i = 1, 2, \ldots, N$ **do**
    Determine the closest cluster by solving

$$k^* = \arg \max_{k=1,2,\ldots,B} \sum_{n \in \mathcal{T}} \mathbb{E}\left[\log_2\left(1 + S_i c_k(i,n)\right)\right]. \quad (20)$$

    Set $\mathcal{C}_{k^*} = \mathcal{C}_{k^*} \cup i$.
  end for
  for $k = 1, 2, \ldots, B$ **do**
    For each cluster, determine optimal single-antenna placement

$$p^* = \max_{p \in \mathcal{C}_k} \sum_{i \in \mathcal{C}_k} \rho_i \sum_{n \in \mathcal{T}} \mathbb{E}\left[\log_2\left(1 + S_i c_p(i,n)\right)\right]. \quad (21)$$

    Set $\phi_{p^*,new} = 1$.
  end for
end while
Set output placement $\vec{\phi}_c = \vec{\phi}_{old}$.

Finally, we study the complexity of Algorithm 3.

Proposition 5. *Algorithm 3 has a complexity of* $\mathcal{O}(BN^2|\mathcal{T}|T_\varepsilon)$.

Proof: Refer Appendix. ∎

We point out that the computational complexity of the clustering approach is independent of available bandwidth at the location. As further motivation, we point out that, in our experiments, the algorithm typically terminates in fewer than ten iterations.

X. Numerical Results on Placement

In this section, we describe the simulation setup for the placement problem, and present the corresponding results.

Simulation setup: We conduct our evaluation of the proposed placement algorithms using the spectrum data shown in Section II. For each of the 8 sq. km regions shown in Figures 4 and 5, we study these algorithms for varying network sizes. Broadly, we consider two different metrics to compare the different placement schemes. As an indicator of general network performance, we compare the *average per-user rates* described in (6) and (8). In addition, we also compare the *percentage of users* in the network who obtain rates above a threshold (2 Mbps or 5 Mbps) as an indicator of the extent of satisfactory coverage that each scheme can provide. In our evaluation, we assume that each base-station is made of an array of $M = 16$ mode-I elements, the size required in Part-I to support 2Mbps per user in a $D = 2$ km cell. [7]

Baseline for comparison: For baseline comparison, we make the assumption that over a given coverage region, users are uniformly distributed. Under this setting, if one were to ignore spatial variation, the optimal scheme would be to uniform placement of network base-stations across the region. We compare the performance of our algorithms against that of this base-line scheme in our experiments. While more elaborate location-specific models on user-density distributions are possible (like those discussed in Part I), this approach allows us to isolate the effect of spatial variability, and examine its impact on the placement problem. Through similar reasoning, we compare our spectrally placed array network against uniformly placed arrays instead of uniformly placed fixed devices.

Figure 13 compares the average per-user rates $R_{MRT/AS}(.)$

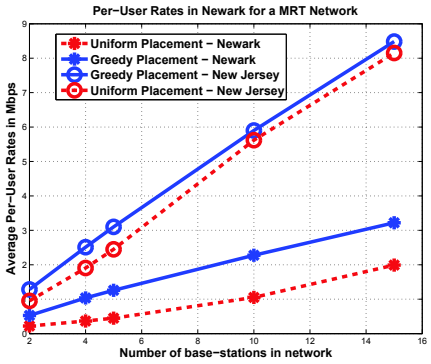

(a) Per-user rates for MRT

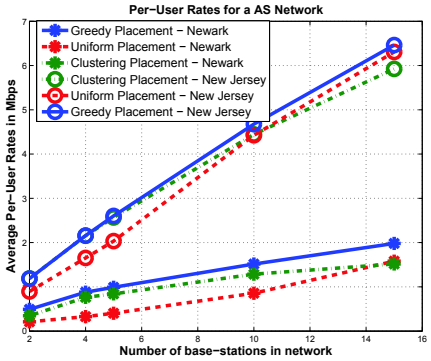

(b) Per-user rates for AS

Fig. 13. Average per-user rates in New Jersey and Newark

achieved by the different algorithms for various network sizes. Let us first consider the MRT case shown in Figure 13(a).

1) In the case of Newark, the rate improvements due to spectrum-aware placement are significant, with the greedy algorithm providing 2.3 times the average rate as uniform placement for the two base-station case (130% improvement), and 1.6 times (60% improvement) for fifteen base-stations . We point out that the convergence between the two schemes with increasing network size is expected since for larger sizes, all placement schemes would converge to a full-coverage scenario.
2) In the case of New Jersey, where the average spectrum availability is significantly higher than Newark (refer Figures 4 and 5), the penalty for poor placement is lesser and consequently, the impact of spectrum-aware placement is lower. Nevertheless, gains are still considerable in this case, with a 35% improvement for a two base-station network and about 5% improvement in average-

[7] Our results in this section extend up to network sizes of 15 base-stations over a 8 sq. km area, implying that on average, each station covers an area of 4 sq. km.

rate for a 10-base-station network.

In summary, spectrum-aware placement gives significant improvement in average rate, especially in spectrum-starved regions. Next, in the case of AS diversity, whose results are presented in Figure 13(b), we make the following observations.

1) The average rates of the AS network are lower than the MRT scheme as expected because of the lack of transmit diversity between the various base-stations.
2) All the trends in the performance of the greedy placement algorithm vs. the uniform placement are consistent with the results from the MRT case, with greedy placement consistently outperforming uniform placement.
3) Interestingly, the clustering-based algorithm steadily under-performs with respect to the greedy algorithm. At lower network sizes, its performance is comparable to that of greedy. However, as network-size grows, its performance deteriorates and approaches that of Uniform placement. This leads us to conclude that the clustering algorithm is more useful for small networks, where it forms a low-complexity alternative to the greedy algorithm.

Next, we study the percentage of users under each scheme who receive rates above a certain threshold rate (2 Mbps, 5 Mbps etc.). This metric, which is an indicator of high-rate coverage, is especially interesting in whitespaces, which is likely to become a vehicle for offload of bandwidth-intensive applications like video streaming. Figures 14 and 15 show these coverage statistics for the MRT and AS networks respectively.

In Newark, we see that the greedy approach considerably outperforms the uniform placement for both MRT and AS networks (see Figures 14(a), 15(a)) for both the 2 Mbps and the 5 Mbps thresholds. Clearly, the greedy placement provides better throughputs to a higher fraction of users in this case. On the other hand, the case of New Jersey is more complex. Though the fraction of users with rates above 2 Mbps is higher initially, the gains vanish at larger network sizes. In the 5 Mbps case, however, the greedy approach clearly supports more users even in the case of New Jersey.

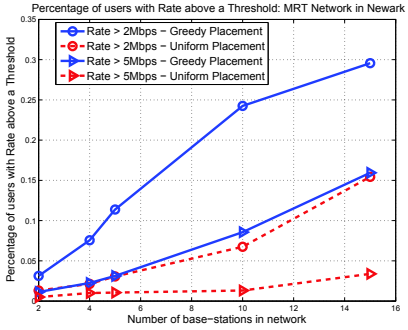

(a) Network in Newark

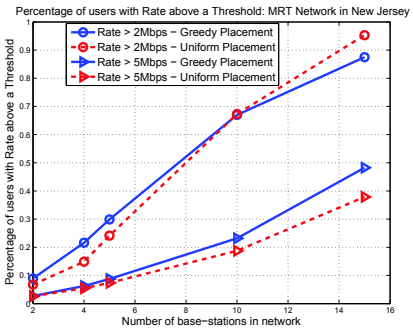

(b) Network in New Jersey

Fig. 14. Percentage of users with rates above a threshold : MRT networks

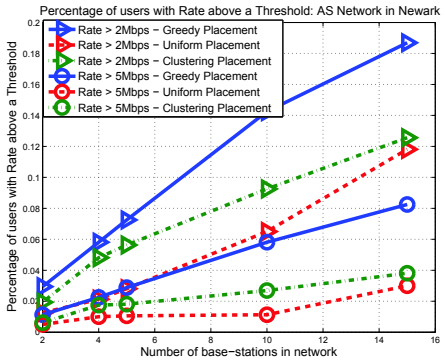

(a) Network in Newark

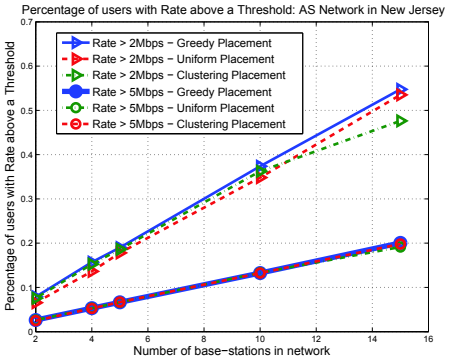

(b) Network in New Jersey

Fig. 15. Percentage of users with rates above a threshold : AS networks

We conclude that in general, the greedy placement approach is able to provide higher peak throughputs to users than the uniform placement algorithm, making it more preferable especially in offload scenarios. As before, we also notice that for AS networks, the clustering algorithm performs close to the greedy algorithm at lower network sizes (upto five), but deteriorates as network size grows.

XI. Future Directions

As mentioned in the introduction, the design of high data-rate systems for TV whitespaces is still at a nascent stage. Our work is a concrete step in this direction, where we make two-fold throughput gains by going from uniformly placed fixed base-stations to spectrally-distributed mode-I arrays. Some problems that are not addressed in this paper include the design of scheduling algorithms under general multi-user MIMO strategies for base-stations. Computational efficiency becomes even more critical in whitespaces since the number of OFDMA sub-bands can be significantly large when many TV bands are available for communication. In the case of base-station placement, exploration of more efficient placement algorithms would be interesting. This third angle of variation would be

the final ingredient in making a network completely tailored to whitespaces.

REFERENCES

[1] H. Ganapathy, M. Madhavan, M. Chetlur and S. Kalyanaraman, "On exploiting degrees-of-freedom in TV whitespaces", *In Proc. IEEE INFOCOM 2012*, Orlando, USA, Mar. 2012.

[2] "FCC frees up vacant TV airwaves for 'Super-Wi-Fi' technologies", Federal Communications Commission, www.fcc.gov.

[3] "Second memorandum opinion and order", Federal Communications Commission, http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db0923/FCC-10-174A1.pdf, Sep. 2010.

[4] "ShowMyWhiteSpace - Locate TV White Space Channels", SpectrumBridge Inc., http://www.spectrumbridge.com

[5] "Telcordia Whitespaces Database", https://prism.telcordia.com/tvws/home

[6] T. L. Marzetta, "Beyond LTE: Hundreds of base station antennas!", *IEEE Communications Theory Workshop 2010*, Cancun, Mexico, May 2010.

[7] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas", *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 3590-3600, Nov. 2010.

[8] D. Gesbert, T. Ekman and N. Christophersen, "Capacity limits of dense palm-sized MIMO arrays", *in Proc. GLOBECOM 2002*, pp. 1187-1191, Taipei, Taiwan, Nov. 2002.

[9] J. Hoydis, S. t. Brink and M. Debbah, "Massive MIMO: How many antennas do we need?", [Online]. Available: http://arxiv.org/abs/1107.1709v1.

[10] R. Murty, R. Chandra, T. Moscibroda and P. Bahl, "SenseLess: A database-driven white spaces network", *In Proc. IEEE DySpan 2011*, Aachen, Germany, May 2011.

[11] C. Cordiero, K. Challapali, D. Birru and S. Shankar, "IEEE 802.22: The first worldwide wireless standard based on cognitive radios", *In Proc. IEEE DySpan 2005*, Nov. 2005.

[12] A. M. Wyglinski, M. Nekovee and Y. T. Hou, "Cognitive radio communications and networks: Principles and practice", Chp. 14, Elsevier, Dec. 2009.

[13] C. Ghosh, S. Roy and D. Cavalcanti, "Coexistence challenges for heterogeneous cognitive wireless networks in TV white spaces", *IEEE Wireless Communication Magazine*, Sep. 2011.

[14] P. Camarda, C. Cormio and C. Passiatore, "An exclusive self-coexistence (ESC) resource sharing algorithm for cognitive 802.22 networks", *In Proc. IEEE Inter. Sym. Wireless Pervasive Computing (ISWPC) 2010*, May 2010.

[15] K. Bian and J.-M. Park, "A coexistence-aware spectrum sharing protocol for 802.22 WRANs", *In Proc. Inter. Conf. on Computer Commun. and Networks*, Washington DC, USA, 2009.

[16] C. Ghosh, D. P. Agrawal, M. B. Rao and C. Cordeiro, "Channel capacity optimization in cooperative cognitive radio networks using game theory", *ACM MC2R Special Issue on Dynamic Spectrum Access and Cognitive Radio*, vol. 13, pp. 2-12, Apr. 2009.

[17] D. Niyato, E. Hossain and Z. Han, "Dynamic spectrum access in IEEE 802.22-based cognitive wireless networks: A game theoretic model for competitive spectrum bidding and pricing", *IEEE Trans. Wireless Commun.*, vol. 16, pp. 16-23, Apr. 2009.

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

[19] "Antenna Structure Registration", http://wireless2.fcc.gov/UlsApp/AsrSearch/asrRegistrationSearch.jsp.

[20] "Find cell tower locations", http://www.cellreception.com/.

[21] K. Johansson, A. Furuskar, P. Karlsson and J. Zander, "Relation between base station characteristics and cost structure in cellular systems", *Proc. of IEEE PIMRC 2004*, pp. 2627-2631, Barcelona, Spain, Sep. 2004.

[22] A. Bogdanov, E. Maneva and S. Riesenfeld, "Power-aware base-station positioning for sensor networks", *In Proc. IEEE INFOCOM 2004*, pp. 575-585, Mar. 2004.

[23] S. S. Dhillon and K. Chakrabarty, "Sensor placement for effective coverage and surveillance in distributed sensor networks", *In Proc. IEEE Wireless Communications and Networking 2003*, vol. 3, pp. 1609-1614, Mar. 2003.

[24] A. Efrat, S. Har-Peled and J. S. B. Mitchell, "Approximation algorithms for two optimal location problems in sensor networks", *In Proc. IEEE BroadNets 2005*, pp. 714- 723, Oct. 2005.

[25] J. Zhang and J. Andrews, "Distributed antenna systems with randomness", *IEEE Trans. Wireless Communications*, vol. 7, pp. 3636-3646, Sep. 2008.

[26] "Techniques and trends in signal monitoring, frequency management and geolocation of wireless emitters", *Aglient whitepaper*.

[27] M. Hata, "Empirical formula for propagation loss in land mobile radio services", *IEEE Trans. Vehicular Technology*, VT-29, pp. 317 - 325, 1980.

[28] Y. Xiao, "WiMAX/MobileFi: Advanced research and technology", Auerbach Publications, ISBN 978-1420043518, 2007.

[29] C. So-In, R. Jain and A.-K. Tamimi, "Capacity evaluation for IEEE 802.16e Mobile WiMAX", *Journal of Computer Systems, Networks, and Communications*, Article ID 279807, Hindawi Publishing Co., 2010.

[30] D. Tse, P. Viswanath, "Fundamentals of wireless communicatoins", Cambridge University Press, 2005.

[31] A. Krause, A. Singh and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies", *Journal of Machine Learning Research*, vol. 9, pp. 235-284, 2008.

[32] L. Wolsey, "An analysis of the greedy algorithm for the submodular set covering problem", *Combinatorica*, vol. 2, pp. 385-393, 1982.

[33] P. R. Goundan and A. S. Schulz, "Revisiting the greedy approach to submodular set function maximization", Jan. 2009.

[34] G. Calinescu, C. Chekuri, M. Pal and J. Vondrak, "Maximizing a submodular set function subject to a matroid constraint (Extended Abstract)", *LNCS, Proc. 12th Intern. Conf. Integer Prog. and Comb. Optimization*, pp. 182 - 196, Ithaca, NY, 2007

[35] J. Vondràk, "Submodularity in combinatorial optimization", *Ph.D. thesis*, Charles University, Prague, 2007.

[36] G. L. Nemhauser and L. A. Wolsey, "Best algorithms for approximating the maximum of a submodular set function", *INFORMS*, vol. 3, pp. 177-188, Aug. 1978.

[37] K. Koh, S. J. Kim and S. Boyd, "An interior-point method for large-scale L1-regularized logistic regression", *Journal of Machine Learning Research*, vol. 8, pp. 1519-1555, Jul. 2007.

[38] A. P. Dempster, N. M. Laird and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm". *Journal of the Royal Statistical Society, Series B*, vol. 39 (1): pp. 138, 1977.

[39] S. P. Lloyd, "Least squares quantization in PCM", *IEEE Trans. Info. Theory*, vol. 28, pp. 129-137, 1982.

[40] K. Nagano, Y. Kawahara and S. Iwata, "Minimum average cost clustering", *in Proc. Neural Info. Proc. Systems*, pp. 1759-1767, Dec. 2010.

APPENDIX

Proof of Proposition 1: The proof follows from two ingredients. First, the rate region $\mathcal{V}_{LT}(n_I)$ is convex for each $\vec{n}$ since our setting has a single centralised controller, thus permitting time-sharing across scheduling policies. Secondly, the service rate achieved by the max-sum-rate scheduler is the solution to $\arg\max_{\nu\in\mathcal{V}_{LT}(n_I)}\sum_k \nu_k$, which essentially corresponds to the "upper-right corner" of the rate region. ■

Proof of Theorem 1: *(i)* Let us define $f_{nlh}(\mathcal{P}) = \log_2\left(1+S_i\sum_{p\in\mathcal{A}_n(\mathcal{P})}\alpha_{ip}|h_{inp}|^2\right)$. From Lemma 1, to show the sub-modularity of $R_{MRT}(\mathcal{P})$, it is sufficient to show that $f_{nlh}(\mathcal{P})$ is sub-modular over $\mathcal{P}$ for any n, l and $\{|h_{inp}|^2\}_{p\in\mathcal{P}}$. Now, consider two placement sets $\mathcal{P}_1$ and $\mathcal{P}_2$ such that $\mathcal{P}_1 \subseteq \mathcal{P}_2 \subseteq \theta$ and consider placing an antenna at location $e \in \theta \setminus (\mathcal{P}_1 \cup \mathcal{P}_2)$. To prove the sub-modularity of f_{nlh}, we need to show that for any i,

$$f_{nlh}(\mathcal{P}_2 \cup e) - f_{nlh}(\mathcal{P}_2) \leq f_{nlh}(\mathcal{P}_1 \cup e) - f_{nlh}(\mathcal{P}_1). \qquad (22)$$

The left-hand-side of the inequality can be computed as

$$\begin{aligned} f_{nlh}(\mathcal{P}_2 \cup e) - f_{nlh}(\mathcal{P}_2) = \quad & \log_2\left(1+S_i\sum_{p\in\mathcal{A}_n(\mathcal{P}_2\cup e)}\alpha_{ip}|h_{inp}|^2\right) \\ & -\log_2\left(1+S_i\sum_{p\in\mathcal{A}_n(\mathcal{P}_2)}\alpha_{ip}|h_{inp}|^2\right) \\ = \quad & \log_2\left(1+\frac{S_i\sum_{p\in\mathcal{A}_n(e)}\alpha_{ip}|h_{inp}|^2}{1+S_i\sum_{p\in\mathcal{A}_n(\mathcal{P}_2)}\alpha_{ip}|h_{inp}|^2}\right), \end{aligned}$$

while the right-hand-side is given by

$$f_{nlh}(\mathcal{P}_1 \cup e) - \quad f_{nlh}(\mathcal{P}_1) = \log_2\left(1+\frac{S_i\sum_{p\in\mathcal{A}_n(e)}\alpha_{ip}|h_{inp}|^2}{1+S_i\sum_{p\in\mathcal{A}_n(\mathcal{P}_1)}\alpha_{ip}|h_{inp}|^2}\right).$$

Since $\sum_{p\in\mathcal{A}_n(\mathcal{P}_1)}\alpha_{ip}|h_{inp}|^2 \leq \sum_{p\in\mathcal{A}_n(\mathcal{P}_2)}\alpha_{ip}|h_{inp}|^2$, the proof of *(i)* follows.

(ii) For the antenna selection case, we set $f_{nlh}(\mathcal{P}) = \log_2\left(1+S_i\max_{p\in\mathcal{A}_n(\mathcal{P})}\alpha_{ip}|h_{inp}|^2\right)$, and begin by making a simple observation. Let

$$\mathcal{X}^*(\mathcal{P}_k) = \arg\max_{p\in\mathcal{A}_n(\mathcal{P}_k\cup e)}\alpha_{ip}|h_{inp}|^2,\ k=1,2.$$

If $e\in\mathcal{X}^*(\mathcal{P}_2)$, then $e\in\mathcal{X}^*(\mathcal{P}_1)$ because $\mathcal{P}_1\subseteq\mathcal{P}_2$. Since $f_{nlh}(\mathcal{P})$ is a non-decreasing function, (22) is straightforward to see when $e\in\mathcal{X}^*(\mathcal{P}_2)$.

On the other hand, when $e\notin\mathcal{X}^*(\mathcal{P}_2)$, then $\mathcal{X}^*(\mathcal{P}_2)\subseteq\mathcal{P}_2 = \mathcal{P}_1\cup(\mathcal{P}_2\setminus\mathcal{P}_1)$. Then $f_{nlh}(\mathcal{P}_2\cup e) - f_{nlh}(\mathcal{P}_2) = 0$ and (22) is trivially satisfied again since $f_{nlh}(\mathcal{P})$ is a non-decreasing function. The result follows. ■

Proof of Proposition 2: The single node rate is given by $R(\{p\}) = \sum_{i=1}^{N}\rho_i\sum_{n\in\mathcal{T}}\Delta_{ipn}\mathbb{E}\left[\log_2\left(1+S_i\alpha_{ip}|h_{inp}|^2\right)\right]$. The inner summation may be simplified to $\sum_{n\in\mathcal{T}}\Delta_{ipn}\mathbb{E}\left[\log_2\left(1+S_i\alpha_{ip}|h_{inp}|^2\right)\right] = |\mathcal{T}(i)\cap\mathcal{T}(p)|\mathbb{E}\left[\log_2\left(1+S_i\alpha_{ip}|h_{inp}|^2\right)\right]$ since fading statistics are assumed identical across TV bands. The result follows. ■

Proof of Proposition 3: *(i)* Since Algorithm 1 adds one node per iteration, it terminates in B iterations. In every iteration, it considers at most $\mathcal{O}(N)$ potential node placements and for each placement, it computes a sum of $\mathcal{O}(N|\mathcal{T}|)$ terms. Each term involves the computation of an effective SINR which costs $\mathcal{O}(W_{\max})$. The result follows.
(ii) Due to the feasibility assumption, the algorithm terminates in at most N iterations. The rest of the arguments may be re-used from Part *(i)*. ■

Proof sketch for Theorem 3: In the low-SINR regime, we may use the well-known approximation $\log(1+x)\approx x$ to re-express (15) as

$$\vec{\phi}^* = \arg\max_{\vec{\phi}\in[0,1]^{|\theta|},\ \sum_l\phi_l\leq B}\vec{\phi}^T\left[\sum_i\rho_iS_i\sum_{n\in\mathcal{T}}\mathbb{E}\left[\vec{c}(i,n)\right]\right]. \quad (23)$$

The result follows since the optimal solution to the linear program in (23) is in fact integral and given by (19). ■

Proof sketch for Proposition 4: First, the function R_{AS} is bounded above by the maximum of the average capacities of the all possible base-station combinations of size B in the considered square cell. Since $\epsilon>0$, and since we solve successive maximization problems in each iteration on this function, each iteration will exhibit diminishing returns. This guarantees termination. ■

Proof of Proposition 5: In each iteration of Algorithm 3, Steps 5-8 have a complexity of $\mathcal{O}(B|\mathcal{T}|)$ operations and Steps 9-12 have a complexity of $\mathcal{O}(B|\mathcal{T}||\mathcal{C}_k|^2) = \mathcal{O}(B|\mathcal{T}|N^2)$ (since $\mathcal{C}_k\leq N$). Consequently, one iteration of Algorithm 3 has a complexity of $\mathcal{O}(NB|\mathcal{T}|+N^2B|\mathcal{T}|) = \mathcal{O}(N^2B|\mathcal{T}|)$. For a total of T_ε iterations, the result follows. ■

K'	Number of users in base-station cell (10 users)
λ'	Rate target for users in bits per second (2Mbps)
D	Size of square cell ($1,2,5$ kms)
N	Square cell is split into N micro-cells
n_I	Number of antenna elements in mode-I array
$\mathcal{T}_d, d\in\{F,I\}$	Spectrum bands available to device d
$P_d, d\in\{F,I\}$	Transmit power limit of device type d (P_F = 4W, P_I = 40 mW)
f_c	Carrier-frequency used for transmission
α_i	Path-loss from user i to base-station array
$h_{kfb}(\mathcal{N})$	Small-scale fading gain from user k to base-station array on OFDM sub-band f of TV band b
$N_{f,int}, N_{i,int}$	Number of fixed and portable interferers respectively
β^F_{kj}	Path-loss from k^{th} user to j^{th} fixed interferer
β^p_{kj}	Path-loss from k^{th} user to j^{th} portable interferer
$\vec{\nu}^{ave}(s,n_I)$	Average rate supported by base-station array of size n_I under scheduling scheme $s\in\Omega$

TABLE I
KEY NOTATION FROM PART-I - TYPICAL VALUES ARE INDICATED IN PARENTHESIS WHERE APPLICABLE.

M	Size of base-station arrays being placed (M = 16)
$\{\rho_i\}$	Spatial distribution of users
D	Size of network coverage region (D = 8 sq. km)
N	Number of micro-cells in coverage region
θ	Set of possible placements $\{1,2,\ldots N\}$ for a single node
α_{ij}	Path-loss from user i to base-station j
h_{inj}	Effective channel gain from user i to base-station j on band n
I_i	Cumulative interference at micro-cell i
N_O	Noise floor
S_i	Power multiplier $\frac{P_I/L}{I_i+N_O}$ for micro-cell i
$\mathcal{P}$	Denotes candidate placement of base-station arrays across region
ϕ_i	Boolean variable with $\phi_i=1$ if base-station array is placed at i
$\mathcal{A}_n(\mathcal{P})$	Subset of $\mathcal{P}$ that has access to band n
$R_{MRT}(\mathcal{P})$	Average network-wide rate when base-stations transmit using MRC
$R_{AS}(\mathcal{P})$	Average network-wide rate when base-stations use antenna selection

TABLE II
KEY NOTATION FROM PART-II - TYPICAL VALUES ARE INDICATED IN PARENTHESIS WHERE APPLICABLE.