PoMeS: Profit-Maximizing Sensor Selection for Crowd-Sensed Spectrum Discovery

Suzan Bayhan[1(✉)], Gürkan Gür[2], and Anatolij Zubow[1]

[1] TU Berlin, Berlin, Germany
{suzan.bayhan,anatolij.zubow}@tu-berlin.de
[2] Zurich University of Applied Sciences (ZHAW), Winterthur, Switzerland
gueu@zhaw.ch

Abstract. In a conventional network management setting, the mobile network operator (MNO) has to account for the traffic fluctuations in its service area and over-provision its network considering the peak traffic. However, this inefficient approach results in a very high cost for the MNO. Alternatively, the MNO can expand its capacity with secondary spectrum discovered opportunistically whenever, wherever needed. While outsourcing the spectrum discovery to a crowd of sensing units may be more advantageous compared to deploying sensing infrastructure itself, the MNO has to offer incentives in the form of payments to the units participating in the sensing campaign. A key challenge for this crowdsensing environment is to decide on how many sensing units to employ given a certain budget under some performance constraints. In this paper, we present a profit-maximizing sensor selection scheme for crowd-sensed spectrum discovery (PoMeS) for MNOs who want to take sensing as a service from the crowd of network elements and pay these sensors for their service. Compared to sensor selection considering the strict sensing accuracy required by the regulations, our heuristics show that an MNO can increase its profit by deciding itself the level of sensing accuracy based on its traffic in each cell site as well as the penalty it has to pay for not satisfying the required sensing accuracy.

Keywords: Spectrum discovery · Crowdsensing · Spectrum sensing

1 Introduction

Mobile network operators (MNO) crave for more radio spectrum to meet the challenging traffic requirements of their customers whose interest is moving towards video-intensive services. Rather than costly over-provisioning, an MNO can expand its capacity with secondary spectrum, which is owned by primary users (PU) but is spatio-temporally unused. However, this opportunistic utilization brings the challenge of discovering the idle spectrum and evicting the

A. Kliks et al. (Eds.): CrownCom 2019, LNICST 291, pp. 3–16, 2019.
https://doi.org/10.1007/978-3-030-25748-4_1

channel when the primary licensed owner appears in the band. While the regulations have moved from spectrum sensing techniques toward spectrum query from white spectrum databases (WSDB) which store all information about the PU transmitters, there is still a need for spectrum sensing as WSDBs aim primarily at protecting the PUs. Moreover, they do not coordinate spectrum sharing among opportunistic secondary users (SU) [2].

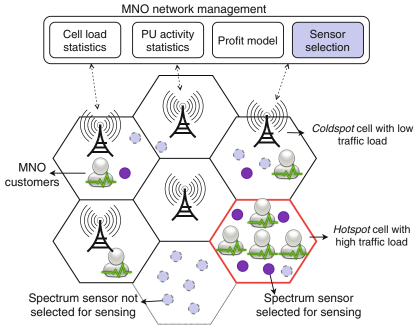

Fig. 1. An MNO can increase its capacity by using the secondary spectrum discovered by the spectrum sensors. The traffic load across the MNO cells (hotspots vs. coldspots) might differ as well as the primary users' activity.

Recent studies [3,5,6,9,12] propose crowdsensing rather than having an MNO deploy its own sensor infrastructure as the latter might lead to higher CAPEX and OPEX. In this paper, we build on previous works, e.g., [3], and address the problem of sensor selection considering the demand as well as the PU traffic activity in each cell. As depicted in Fig. 1, the MNO via its management system can collect statistics about its own traffic as well as the traffic in the PU spectrum to make more informed decision for sensor selection. While deciding on which sensors to select in each cell, the MNO has to consider the related costs of sensing, e.g., payments to the sensors, and its gain via the discovered spectrum. This gain depends on the spectrum bandwidth that will be discovered via sensors and the additional traffic that could be served with that resource. As traffic characteristics may vary among cells, particularly for small cells where the statistical multiplexing is minimal compared to macrocells, deploying sensors may not payoff if the traffic load is low in a cell. While sensing cost and efficiency are important for the MNO's profit, the regulatory bodies assert sensing accuracy requirement to protect the incumbent users, i.e., PUs, which might lead to higher sensing cost for the MNO. For example, requiring a PU detection accuracy as well as the false alarm probability below a certain threshold regardless of the PU traffic or secondary network's traffic might result in wasteful sensing by the sensors. To go beyond this inflexibility and address more realistic scenarios in this work, we explore a relaxed case where the MNO can prefer maintaining a lower sensing accuracy and then pay for the resulting PU collisions due to its lower PU detection performance. This approach provides the MNO flexibility to maintain its sensing accuracy at different levels in each cell depending on the expected PU traffic (and thereby collision-related penalties) and its traffic load (i.e., operation at different points of cost-benefit trade-off).

In this work, we make the following key contributions regarding spectrum discovery and sensor selection for crowdsensing:

- We propose spectrum discovery via crowd-sensing under a budget constraint. While there is a rich literature on crowdsourced sensing and sensor selection such as [3,6,9,12], the business aspects of this problem is largely overlooked. Our goal in this paper is to analyze how the profit model of an MNO might affect its decisions for sensor selection.
- We propose to relax the sensing accuracy constraints to save from the sensing cost, especially for cells without a high traffic demand, and yet motivate the MNO to attain a higher PU detection accuracy. We achieve this goal by introducing a penalty to the MNO if it cannot satisfy the required minimum detection accuracy. However, we believe that there should be no penalty for the divergence from the maximum false alarm probability since the MNO will inherently have incentives to minimize the false alarm probability under high traffic load. In line with that expected behavior, our simulations show that the MNO favors higher PU detection accuracy and low false alarm rate under high traffic load.
- We present a thorough analysis of our proposals for different system parameters including allocated budget, traffic load, and fraction of hot spots. Our simulations show that an MNO benefits from relaxing the strict requirements on the sensing accuracy. In our proposal, an MNO can target different accuracy levels (e.g., lower false alarm if the need for the required capacity is higher) depending on its traffic in each cell site as well as the penalty it has to pay for not satisfying the required sensing accuracy.

The rest of the paper is organized as follows. First, Sect. 2 presents the considered system model for the network and the sensors. Next, Sect. 3 introduces the cost and utility of spectrum discovery. It also formulates the profit-maximizing sensor selection (PoMeS) problem, while Sect. 4 provides several polynomial-time complexity heuristics for PoMeS. Section 5 presents a detailed assessment of the performance of the devised heuristics in comparison to the baseline which has to ensure the sensing accuracy requirements imposed by the regulatory bodies. Section 6 provides an overview of the related work on sensor selection for crowdsourced spectrum discovery while Sect. 7 concludes the paper with a brief discussion of future directions.

2 System Model

Consider an MNO with $\mathcal{A} = [A_i, \cdots, A_K]$ cell sites. Each cell site hosts users with a certain demand denoted by r_i requests/sec. Each request requires a minimum rate, e.g., $c_{\min}$ bits/sec, and for each request the user pays μ monetary units. The bandwidth of the PU's channel is denoted by B Hertz. In this channel, the PU has an activity with probability p_1^i in A_i. Hence, the MNO can use the channel with probability $p_0^i = 1 - p_1^i$ if it can discover the spectrum opportunities with perfect accuracy and without any overhead. However, depending on the length of the sensing and reporting duration as well as number of sensors, there will be an overhead as illustrated in Fig. 2. Moreover, the spectrum sensors might falsely conclude the state of the channel as occupied due to errors in sensing or fluctuations in the channel, which then decreases the amount of discovered spectrum.

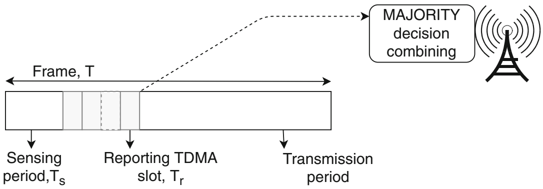

Fig. 2. A frame starts with the sensing period and continues with the reporting period. Each sensor reports its sensing outcome via a TDMA uplink during the reporting slot allocated to it. After completion of the reporting period, the BS at each cell applies majority logic to decide on the spectrum state.

The MNO has a budget of $\mathcal{B}$ (in currency units C) to pay for its crowdsensing campaign. With this budget, it needs to decide on how many sensors (and in case of heterogeneous sensors, considering their sensing capabilities and price of sensing) to employ for sensing in each cell site A_i. We denote the total number of sensors by N and the price of sensing by μ_s C per bps. After the sensors are selected (e.g., S out of N), they start sensing with the requested sensing rate (β_s) during the sensing period T_s in a frame with duration T. Each sensor reports its one bit sensing outcome, i.e., {0: idle, 1: busy}, to the base station (BS) of its cell using TDMA in a slot of duration T_r as in [8]. As a result, the time left after sensing and in-band reporting in a frame of length T equals to $T - T_s - ST_r$. Hence, the normalized sensing overhead is $\omega = \frac{T_s + ST_r}{T}$. We assume identical sensing characteristics for the sensors, i.e., they have identical local probability of false alarm (P_f) and local probability of PU detection (P_d). Given that each sensor applies energy detection, we can derive the P_f and P_d values based on SNR and noise at each sensor [1,8].

After collecting the sensing outcomes, the BS fuses this sensory data (H_i denoting sensor i's binary decision) from S sensors using majority logic, which is known to be robust against sensing errors [3]. Simply, the BS checks if $\sum_{i=1}^{S} H_i \geqslant \lceil S/2 \rceil$. If this inequality holds, it concludes that the spectrum is in use by its primary owners, hence it cannot use this spectrum. Otherwise, it can serve its users through this spectrum, e.g., via carrier aggregation with the existing licensed spectrum as currently specified in LTE. We denote the spectral efficiency of the MNO by κ bps/Hz.

Since regulatory bodies target high utilization of this scarce resource without harming the PUs, they might assert certain sensing accuracy constraints at a cell: Q_f denoting global probability of false alarm and Q_d denoting global probability of PU detection in a cell. We assume that the MNO is required to sustain Q_d^* and can target different Q_f based on the user demand in each cell site. However, for cell sites where the MNO has only low traffic activity (yet higher than the available capacity), it might consider employing a lower number of sensors which would result in lower Q_d, possibly lower than Q_d^*. In this case, we assume that the MNO will have to pay a certain penalty for not meeting the required accuracy. This approach aims at relaxing the strict requirements on the detection accuracy and giving the MNO ability to maintain different sensing accuracy levels across its cells. Yet, the MNO would essentially be driven towards higher sensing accuracy due to the penalty mechanism. In the next section, we introduce our proposal for sensor selection in a wireless network represented by this system model.

3 Profit-Maximizing Sensor Selection (PoMeS)

Let us first define the utility of crowdsensing in terms of the amount of spectrum discovered by the sensors. If there are m sensors participating in sensing, then the amount of the spectrum that will be discovered can be calculated as:

$$\mathcal{U}(m) = p_0 B \left(\frac{T - T_s - mT_r}{T} \right) (1 - Q_f(m)) \text{ Hz}, \tag{1}$$

where $Q_f(m)$ is the false alarm probability if m sensors participate in sensing. We can calculate false alarm probability for majority voting as below [3,8]:

$$Q_f(m) = \sum_{n=\lceil \frac{m}{2} \rceil}^{m} \binom{m}{n} (P_f)^n (1 - P_f)^{m-n}. \tag{2}$$

Similarly, we calculate $Q_d(m)$ as follows:

$$Q_d(m) = \sum_{n=\lceil \frac{m}{2} \rceil}^{m} \binom{m}{n} (P_d)^n (1 - P_d)^{m-n}. \tag{3}$$

We can model the profit of an MNO from each cell considering the number of requests that will be served with the discovered capacity. The discovered capacity is simply $\mathcal{U}_i \kappa$ bps. Hence, the number of requests that can be served with this capacity is:

$$R_i^{\max} = \min(r_i, \frac{\mathcal{U}_i \kappa}{c_{\min}}) \text{ requests/s}. \tag{4}$$

Consequently, we calculate the MNO's income in currency-per-second (C/s) from its customers in A_i as follows:

$$\Pi_i^+ = R_i^{\max} \mu \quad \text{C/s}. \tag{5}$$

Moreover, the MNO has to pay for the sensing service to the selected N_i sensors. If each sensor has to perform sensing with rate β_s bps and the cost of sensing is μ_s per sensing bit, then the total payment for A_i is as follows:

$$\Pi_i^- = \mu_s \beta_s N_i \text{ C/s.} \tag{6}$$

We introduce a penalty of low PU detection accuracy for the MNO to avoid low sensing accuracy. We denote this penalty by μ_c and the gap between the required Q_d and the realized one by $\Delta Q_{d,i} = \max(0, Q_d^* - Q_{d,i})$ for cell A_i. The resulting penalty in monetary terms equals to $\mu_c \Delta Q_{d,i} R_i^{\max}$. Note that other options for the penalty function is possible, e.g., an exponential function of $\Delta Q_{d,i}$ which is more punishing compared to the linear function used here. Next, we calculate the net profit using (5) and (6), which gives us the following:

$$\Pi_i = R_i^{\max}\mu - \mu_s \beta_s N_i - \mu_c R_i^{\max} \max(0, Q_d^* - Q_{d,i}). \tag{7}$$

Now, let us present the optimization problem which will be solved by the MNO to decide on the number of sensors to be selected for each cell. Profit-maximizing sensor selection (PoMeS) problem is formally defined as follows:

$$\max_{N_i} \sum_{A_i \in \mathcal{A}} R_i^{\max}\mu - \mu_s \beta_s N_i - \mu_c R_i^{\max} \max(0, Q_d^* - Q_{d,i}) \tag{8}$$

subject to:

$$\mu_s \beta_s (\sum N_i) \leqslant \mathcal{B} \tag{9}$$

$$N_i \leqslant \lfloor \frac{T - T_s}{T_r} \rfloor \quad \forall A_i \in \mathcal{A} \tag{10}$$

$$N_i \geqslant 0 \quad \forall A_i \in \mathcal{A}. \tag{11}$$

Constraints (9) determines the maximum number of sensors to employ in the sensing campaign due to the budget constraint while Const. (10) restricts the number of sensors due to the finite size of the frame. While the constraints are linear in the decision variable N_i, the objective function is non-linear function due to non-linearity of (1), (2), and (3). Hence, our problem is a non-linear integer problem whose complexity is typically high. In addition, our problem has to account for combinations across all cells, which makes this problem computationally hard. Hence, we devise polynomial complexity heuristics in the next section.

4 Sensor Selection Heuristics

Equal Budget per Cell (EQ): This heuristic divides the budget by the number of total cells and finds the best decision for each cell under the cell's budget constraint, i.e., $\mathcal{B}/K$. Then, the maximum number of sensors that could be employed is $N_{\max} = \min(\lfloor \frac{T-T_s}{T_r} \rfloor, \lfloor \frac{\mathcal{B}}{K \mu_s \beta_s} \rfloor)$. Next, this heuristic exhaustively searches for the setting that maximizes the objective function Π_i with constraint

$N_i \leqslant N_{\max}$. If $\Pi_i^{\max}$ is nonnegative, the corresponding number of sensors is assigned to this cell. Otherwise, no sensor is deployed. Note that if the sensors have different cost and sensing accuracy, then the sensor selection problem would be more complex. Here, due to our assumption of a homogeneous setting, EQ only needs to decide on the number of sensors. Moreover, although the MNO's profit is zero for the considered time period, i.e., $\Pi_i^{\max} = 0$, the MNO may still prefer deploying sensors to this cell to increase its service availability. Because better availability might increase the reputation of the MNO, which may attract more customers in the long run. In case $N_{\max}$ is zero which is expected to happen under low budget, EQ finds the minimum required number of sensors for satisfying Q_d^*. Then, EQ selects the cells with the highest loads and employs the minimum number of sensors satisfying Q_d^* in those cells. Computational complexity of EQ is $\mathcal{O}(KN_{\max})$ as EQ calculates the profit for each cell while considering each possible number of sensors lower than or equal to $N_{\max}$.

Budget Proportional to the Serving Capacity of the cell (PROP): Rather than allocating equal budget to each cell, this heuristic allocates the budget proportional to the number of requests $R_i^{\max}$ that could be served by each cell. For the ease of computation, we set $Q_{f,i} = 0$. Then, the maximum number of requests that could be served for a cell equals to

$$R_i^{\max} = \min(r_i, \frac{Bp_0^i(1-\omega)\kappa}{c_{\min}}).$$

Then, PROP allocates to cell A_i a budget of $\mathcal{B}_i$:

$$\mathcal{B}_i = \frac{\mathcal{B}R_i^{\max}}{\sum_{A_i \in \mathcal{A}} R_i^{\max}}. \tag{12}$$

Under $\mathcal{B}_i$, PROP exhaustively searches for the best number of sensors to be selected for each cell. Computational complexity of PROP is $\mathcal{O}(KN_{\max})$.

Incremental Gain Based Greedy Assignment (INGA): In this case, the sensor allocation starts from the cell with the highest incremental gain defined as $\Delta\Pi_i = \Pi_{i,m+1} - \Pi_{i,m}$ where m is the current number of sensors deployed in the cell. One more sensor is allocated to this cell with the highest $\Delta\Pi_i$ resulting in $m+1$ sensors in the cell. The iteration continues with the next cell which attains the highest incremental gain with one more sensor deployed. The assignment halts either when globally budget is depleted or when maximum $\Delta\Pi_i$ is negative. Complexity of INGA is $\mathcal{O}(NK\log(K))$ as INGA finds for each sensor the cell with the maximum incremental gain via a sorting algorithm.

Baseline Satisfying (Q_d^*, Q_f^*) Required by the Regulatory Body (REG): This heuristic has two variants: REG-EQ and REG-PROP where the former follows the same budget allocation approach as EQ and the latter as PROP. However, while performing exhaustive search, a solution is considered to be feasible only if the sensing constraints for both Q_d^* and Q_f^* are satisfied. If the cell budget is not sufficient to deploy the minimum required number of sensors, then

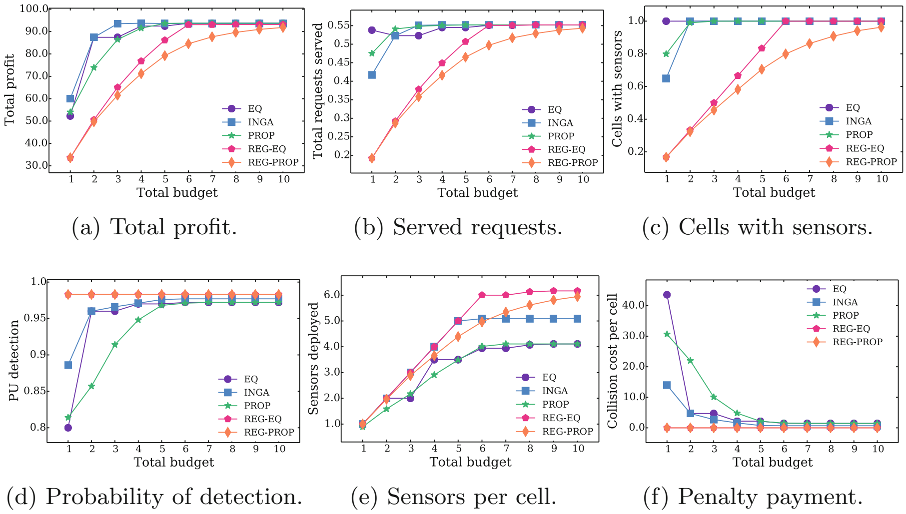

(a) Total profit. (b) Served requests. (c) Cells with sensors.

(d) Probability of detection. (e) Sensors per cell. (f) Penalty payment.

Fig. 3. Impact of increasing budget $\mathcal{B}$ allocated for the spectrum sensing service by the MNO.

sensing service is not available for this cell resulting in no additional spectrum in the cell. We consider this scheme to be the baseline which regulatory bodies have proposed in earlier proposals, e.g., (0.9, 0.1) for IEEE 802.22 [8]. Computational complexity of REG is the same as EQ and PROP, i.e., $\mathcal{O}(KN_{\max})$.

5 Performance Evaluation

Here, after describing the simulation setting and the selected parameters, we present the numerical results where we investigate the impact of various system parameters, e.g., budget, on the performance.

5.1 Simulation Setting

We simulate a cellular network with $K = 2000$ cell sites using our custom Python simulator and analyze the performance of our proposed schemes. The PU's off probability is distributed uniformly in [0.2, 0.8]. Other parameters (listed in Table 1) are set as follows: $\mu_s = 1$, $\mu = 1$, $\mu_c = 5$, $\kappa = 10$ bps/Hz, $(P_d, P_f) = (0.8, 0.1)$, and $(Q_d^*, Q_f^*) = (0.98, 0.05)$. For generating the request distribution, we first pick randomly σ of the cells as hotspots. Total requests generated from these cells will account for R_σ fraction of the requests. The rest which we call as *coldspots* will account for $(1 - R_\sigma)$ fraction of the requests. In each cell category, to have some variance in traffic, we generate the requests uniformly distributed in an interval, e.g., with 10% variance from the average load. If not stated otherwise, we set $R_\sigma = 0.6$ and $\sigma = 0.1$.

Table 1. Key simulation parameters.

Parameter	Value
Number of cell sites (K)	300
Av. num. of requests per cell (R_i)	90
Duration of frame, sensing period, and reporting slots (T, T_s, T_r)	(10, 1/6*10^{-3}, 10*10^{-3}) ms
PU's idle probability (p_0^i)	U(0.2, 0.8)
Sensing price (μ_s)	1
Service price (μ)	1
Collision penalty (μ_c)	5
Min. capacity requirement ($c_{\min}$)	3 Mbps
Spectral efficiency (κ)	10 bps/Hz
Sensor sensing accuracy (P_d, P_f)	(0.8, 0.1)
Target sensing accuracy (Q_d^*, Q_f^*)	(0.98, 0.05)

5.2 Impact of Budget $\mathcal{B}$

In Fig. 3, we increase the total budget $\mathcal{B}$ from 1 to 10 sensors/cell with a step size of one sensor/cell. For example, when $\mathcal{B} = 1$, this means that the MNO's budget in total is $\mu_s \beta_s K$ and therefore can afford only allocating on average one sensor to each cell.

Figure 3a shows the change in the total profit of the MNO. As expected, increasing budget increases the profit. However, each scheme experiences saturation after a certain budget. This is due to the diminishing returns: deploying more sensors only increases the capacity marginally. Hence, the resulting increase in the MNO income via serving more traffic does not justify the increased payment for sensors. Another reason might be that the discovered capacity is already sufficient to serve all traffic in the cell, which invalidates more capacity addition. Observing Fig. 3b, we see that the saturation is due to the first reason as only a maximum of approximately 60% of the requests are served. We observe diminishing returns in sensing accuracy and thereby related utility with increasing number of sensors. Therefore, our schemes might prefer deploying lower number of sensors than the one required by the regulation-conforming heuristics which have to ensure (Q_d^* and Q_f^*). This insight is supported by Fig. 3d and e which show lower Q_d and lower number of sensors deployed for our heuristics, respectively. Figure 3f shows that under low budget our heuristics might sacrifice from sensing accuracy and pay for resulting penalty. In return, more cells can benefit from capacity expansion via opportunistic spectrum discovery. For example, as Fig. 3c shows, our heuristics employ sensors from a wide range of cells in contrast to REG-EQ and REG-PROP schemes. As an example, for $\mathcal{B} = 1$, EQ employs one sensor/cell covering all cell sites while REG-EQ employs sensors only in 17% of the cell sites as the minimum required number of sensors is six in that setting.

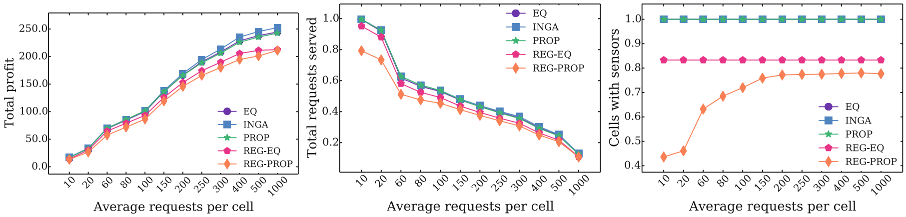

(a) Total profit. (b) Served requests. (c) Cells with sensors.

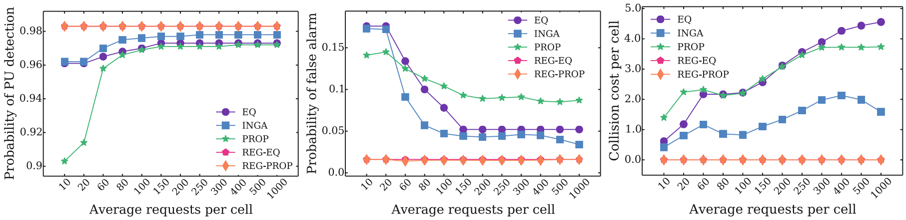

(d) Probability of PU detection. (e) Probability of false alarm. (f) Penalty.

Fig. 4. Impact of increasing load (i.e., number of requests generated per cell).

Comparing all schemes in Fig. 3a, we see that when budget is low (e.g., $\mathcal{B} = 1$), INGA has the highest performance followed by PROP. Later with increasing $\mathcal{B}$, EQ gradually attains similar performance. All these schemes have a significant performance improvement (e.g., reaching 0.8x for low budget) over the regulations-conforming heuristics REG-EQ and REG-PROP. This improvement comes with a trade-off in sensing accuracy as observed in Fig. 3d where REG-EQ and REG-PROP have the highest and exactly the same performance.

5.3 Impact of Traffic Load R_i

In Fig. 4, we plot the impact of increasing load in terms of number of requests per cell. Here, we set $\mathcal{B} = 5$ sensors/cell. Under all loads, INGA succeeds the highest profit with a gap of around 5% from its closest follower PROP. All schemes have higher profit with increasing load, but there seems to be a saturation point after which the profit improves very slightly. The saturation point is reached when all the discovered bandwidth is used for the requests.

Another observation is that if the MNO prefers or has to use REG approaches, it should choose equal division of its budget across its cells as consistently REG-EQ over-performs REG-PROP for about (5–15%) depending on the setting. While increasing load improves the profit of the operator due to higher payments from the customers, some of the requests may be blocked as we see lower values of served requests in Fig. 4b. The selection of which cells to serve differs from one scheme to another. For example, schemes with proportional budget assignment prioritize hotspots, e.g., cells with higher load. Figure 4c shows that

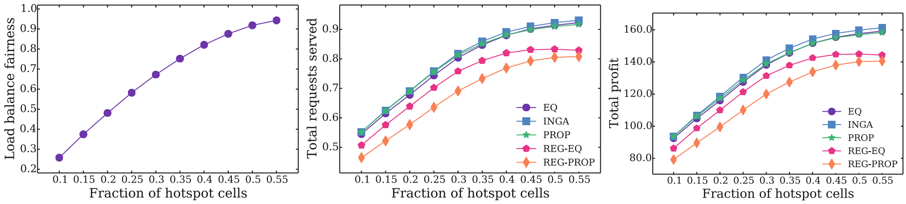

(a) Traffic load balance. (b) Served requests. (c) Total profit.

Fig. 5. Impact of increasing fraction of hotspots.

REG schemes has a confined cell span, i.e., employs sensors only a subset of cells. This is due to the limited budget which can afford 5 sensors per cell, lower than 6 sensors required for (0.98, 0.05) sensing accuracy. Hence, around 5/6($\sim$80%) of the cells are selected for sensor deployment under REG-EQ. For REG-PROP, the fraction of cells with capacity extension is lower as proportional assignment of the budget might employ more sensors in a highly-loaded cell, thus leaving other cells without any sensors.

When it comes to sensing accuracy, we observe in Fig. 4d and e that increasing load requires more spectrum, i.e., lower false alarm probability. Hence, each scheme prefers deploying more sensors, which consequently also improves probability of PU detection. Additionally, since penalty function incurs also the number of requests as its multiplier, the penalty of a lower detection accuracy becomes higher under high load. As a result, MNO is motivated to satisfy minimum Q_d^*. As Fig. 4f shows, INGA results in a lower penalty cost for possible collisions with the PU compared to other heuristics.

5.4 Impact of Fraction of Hotspots σ

In the earlier scenarios, only $\sigma = 0.1$ fraction of the cells are hotspots accounting for 60% of the network traffic. Now, we analyze the impact of traffic heterogeneity across cells. With increasing values of $\sigma = [0.10, 0.55]$, MNO's traffic becomes more homogeneously distributed across cells while the total number of requests remains the same. Figure 5a plots the observed traffic load balance which is calculated as the Jain's fairness index using the distribution of number of requests across all cell sites. As we see in the figure, MNO's traffic balance is almost 0.9 when $\sigma = 0.55$ and around 0.3 when $\sigma = 0.1$.

Figure 5b plots the change in the fraction of served requests with increasing σ. Under a more evenly distributed traffic, all schemes can serve a higher number of requests (Fig. 5b) by discovering additional spectrum in more cells. As a result, the total profit increases as shown in Fig. 5c. The relative performance of each scheme follows the same trend as observed in other scenarios: our proposals outperform REG variants under all traffic heterogeneity settings.

5.5 Discussion and Practical Considerations

Here, we discuss more on the applicability of our heuristics in a practical setting. Our heuristics use traffic demand in each cell site, the PU's channel availability information, and the sensor accuracies. The first two statistics can be collected at each cell site over an observation period. The traffic demand model is easy to acquire using the received requests from the MNO's customers at each BS. As the earlier research shows, e.g., [10], the network traffic has seasonality and is predictable to a good accuracy. For PU channel availability statistics, collected sensor data can be used to derive the PU's traffic pattern and thereby p_0. The prior work on PU traffic prediction, e.g., [7], can be used not only for estimating p_0 but also for more accurate estimation of the PU dynamics. Lastly, each sensor's accuracy can be computed at each BS by comparing it with the final sensing outcome. Obviously, it will take some time to converge to the sensor's accuracy level.

When it comes to executing payment to sensors, an MNO can prefer using smart-contract based solutions as proposed in [3]. Smart contracts are digital counterparts of traditional contracts which define the terms of an agreement as well as dispute resolution approach. However, smart contracts do not need the trust between trading parties as the contract itself is the point of trust. Hence, it can fit to our setting where the sensors and the MNO do not have to trust each other. Nevertheless, using smart contracts might entail additional monetary cost which must be paid eventually by the MNO. Hence, the MNO has to revise its profit calculation based on the additional cost of using a smart contract network.

6 Related Work

The most relevant works to ours elaborate on sensor selection for crowdsourced spectrum sensing, e.g. [3–6,9,11,12]. In [11], Ying et al. design a pricing mechanism with joint consideration of sampling value, data quality and cost of incentivized sensing. Their main contribution resides on the integration of device heterogeneity reflected in sensing data quality and costs. However, they do not consider constraints such as regulatory requirements but focus on the intricacies of REM construction under heterogeneous sensor settings. In [6], Jin et al. similarly elaborate on the participant selection in crowdsourced spectrum sensing systems and model it as a reverse auction problem. Their main focus is the privacy problem in such a system. To this end, they develop a framework for the MNO to select sensing participants in a differentially privacy-preserving manner. However, they do not address sensor selection problem under some regulatory constraints as we do in our work but rather assume that the MNO has pre-determined the sensing locations of each sensing task according to existing methods. In [12], Zhang et al. address the security of crowdsourcing-based spectrum sensing since the cooperative process is vulnerable to malicious sensing data injection attacks. Their approach considers the instantaneous trustworthiness of mobile detectors in combination with their reputation scores during data fusion for sensing decisions. However, their key concern is security rather than optimal sensor selection in compliance with regulatory requirements or profit maximization for MNOs.

There are also works which investigate the practical implications of crowd-sourced spectrum sensing. In [9], Nika et al. propose real-time spectrum monitoring with strong coverage using low-cost and commodity hardware. Although they do not work on sensor selection or sensing constraints due to regulations, their work is especially interesting as a feasibility study and employs practical hardware for large-scale low-cost sensing. In [5], Chakraborty et al. also crowd-source spectrum monitoring to low-cost and low-power commodity devices. To address the overhead drawback in the crowdsourced spectrum sensing, they propose three heuristics to select the minimum number of spectrum sensors that can best estimate the spectrum at the requested locations. In [4], they further develop a crowdsensing framework for low-cost and large-scale settings, which includes a technique for the sensor selection and fusion problem based on sensor decorrelation and clustering. However, they do not consider how the users are incentivized or the broader economical aspects.

Overall, none of these works except [3] consider the operator's business strategy and more specifically its profit. Spass [3] has a similarity with PoMeS and it also aims at maximizing MNO's profit. But, it considers a single cell. Moreover, the MNO considers the monetary overhead due to usage of a smart-contract network to pay the sensors participating in its sensing campaign in that work. Different than Spass, in this paper, we consider multi-cell setting with heterogeneous cell traffic and focus on how many sensors to select in each cell.

7 Conclusion

Opportunistic spectrum access provides ample opportunities for MNOs to increase their capacity cost-effectively whenever, wherever needed. For spectrum discovery, an MNO can buy spectrum sensing service from the sensing-capable sensors in exchange of payments for that service. However, the optimal number of sensors to be employed depends on various factors such as the MNO's own traffic in each cell, spectrum occupancy, and sensing price. In this work, we have first formulated a profit-maximizing sensor selection problem as an integer non-linear problem and then devised several heuristics with polynomial time complexity to decide on how many sensors to select in each MNO cell. Our solutions which might target a lower sensing accuracy at the expense of some monetary penalty outperform traditional approaches which have to ensure a certain level of sensing accuracy regardless of the network dynamics, e.g., needed spectrum, PU traffic, MNO traffic. Moreover, our simulations show that the MNO prefers maintaining lower sensing accuracy only when the network load is low and budget for spectrum sensing payment is limited. As future work, we plan to consider a more heterogeneous setting wherein sensors might be diverse in their sensing accuracy as well as their sensing price.

Acknowledgments. This work was partially supported by the Scientific and Technical Research Council of Turkey (TUBITAK) under grant number 116E245 and by the European Horizon 2020 Programme under grant agreement n688116 (eWINE project).

References

1. Axell, E., Leus, G., Larsson, E.G., Poor, H.V.: Spectrum sensing for cognitive radio: State-of-the-art and recent advances. IEEE Signal Process. Mag. **29**(3), 101–116 (2012)
2. Baig, G., Alistarh, D., Karagiannis, T., Radunovic, B., Balkwill, M., Qiu, L.: Towards unlicensed cellular networks in TV white spaces. In: ACM CONEXT (2017)
3. Bayhan, S., Zubow, A., Wolisz, A.: Spass: spectrum sensing as a service via smart contracts. In: IEEE DySPAN (2018)
4. Chakraborty, A., Bhattacharya, A., Kamal, S., Das, S.R., Gupta, H., Djuric, P.M.: Spectrum patrolling with crowdsourced spectrum sensors. In: IEEE Conference on Computer Communications (INFOCOM) (2018)
5. Chakraborty, A., Rahman, M.S., Gupta, H., Das, S.R.: Specsense: crowdsensing for efficient querying of spectrum occupancy. In: IEEE Conference on Computer Communications (INFOCOM) (2017)
6. Jin, X., Zhang, Y.: Privacy-preserving crowdsourced spectrum sensing. IEEE/ACM Trans. Netw. **26**(3), 1236–1249 (2018)
7. Li, X., Zekavat, S.A.R.: Traffic pattern prediction based spectrum sharing for cognitive radios. In: Cognitive Radio Systems. InTech (2009)
8. Maleki, S., Chepuri, S.P., Leus, G.: Energy and throughput efficient strategies for cooperative spectrum sensing in CRs. In: IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC) (2011)
9. Nika, A., Zhang, Z., Zhou, X., Zhao, B.Y., Zheng, H.: Towards commoditized real-time spectrum monitoring. In: ACM Workshop on Hot Topics in Wireless, HotWireless 2014, pp. 25–30 (2014)
10. Xu, F., Li, Y., Wang, H., Zhang, P., Jin, D.: Understanding mobile traffic patterns of large scale cellular towers in urban environment. IEEE/ACM Trans. Netw. (TON) **25**(2), 1147–1161 (2017)
11. Ying, X., Roy, S., Poovendran, R.: Pricing mechanism for quality-based radio mapping via crowdsourcing. In: IEEE Global Communications Conference (GLOBECOM) (2016)
12. Zhang, R., Zhang, J., Zhang, Y., Zhang, C.: Secure crowdsourcing-based cooperative spectrum sensing. In: IEEE INFOCOM (2013)