Interference in Cellular Networks: The Minimum Membership Set Cover Problem

Fabian Kuhn[1], Pascal von Rickenbach[1], Roger Wattenhofer[1],
Emo Welzl[2], and Aaron Zollinger[1]
kuhn@tik.ee.ethz.ch, pascalv@tik.ee.ethz.ch, wattenhofer@tik.ee.ethz.ch,
welzl@inf.ethz.ch, and zollinger@tik.ee.ethz.ch

[1] Computer Engineering and Networks Laboratory, ETH Zurich, 8092 Zurich, Switzerland
[2] Department of Computer Science, ETH Zurich, 8092 Zurich, Switzerland

Abstract. The infrastructure for mobile distributed tasks is often formed by cellular networks. One of the major issues in such networks is interference. In this paper we tackle interference reduction by suitable assignment of transmission power levels to base stations. This task is formalized introducing the *Minimum Membership Set Cover* combinatorial optimization problem. On the one hand we prove that in polynomial time the optimal solution of the problem cannot be approximated more closely than with a factor $\ln n$. On the other hand we present an algorithm exploiting linear programming relaxation techniques which asymptotically matches this lower bound.

1 Introduction

Cellular networks are heterogeneous networks consisting of two different types of nodes: base stations and clients. The base stations—acting as servers—are interconnected by an external fixed backbone network; clients are connected via radio links to base stations. The totality of the base stations forms the infrastructure for distributed applications running on the clients, the most prominent of which probably being mobile telephony. Cellular networks can however more broadly be considered a type of infrastructure for mobile distributed tasks in general.

Since communication over the wireless links takes place in a shared medium, interference can occur at a client if it is within transmission range of more than one base station. In order to prevent such collisions, coordination among the conflicting base stations is required. Commonly this problem is solved by segmenting the available frequency spectrum into channels to be assigned to the base stations in such a way as to prevent interference, in particular such that no two base stations with overlapping transmission range use the same channel.

In this paper we assume a different approach to interference reduction. The basis of our analysis is formed by the observation that interference effects occurring at a client depend on the number of base stations by whose transmission ranges it is covered. In particular for solutions using frequency division multiplexing as described above, the number of base stations covering a client is a lower bound for the number of channels required to avoid conflicts; a reduction in the required number of channels, in turn, can be exploited to broaden the frequency segments and consequently to increase communication bandwidth. On the other hand, also with systems using code division multiplexing,

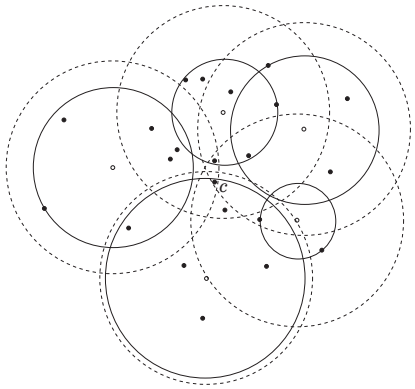

Fig. 1. If the base stations (hollow points) are assigned identical transmission power levels (dashed circles), client c experiences high interference, since it is covered by all base stations. Interference can be reduced by assigning appropriate power values (solid circles), such that all clients are covered by at most two base stations.

the coding overhead can be reduced if only a small number of base stations cover a client.

The transmission range of a base station—and consequently the coverage properties of the clients—depends on its position, obstacles hindering the propagation of electromagnetic waves, such as walls, buildings, or mountains, and the base station transmission power. Since due to legal or architectural constraints the former two factors are generally difficult to control, we assume a scenario in which the base station positions are fixed, each base station can however adjust its transmission power. The problem of minimizing interference then consists in assigning every base station a transmission power level such that the number of base stations covering any node is minimal (cf. Figure 1). At the same time however, it has to be guaranteed that every client is covered by at least one base station in order to maintain availability of the network.

In Figure 1 the area covered by a base station b transmitting with a given power level is represented by a disk centered at b and having a radius corresponding to the chosen transmission power. Practical measurements however show that this idealization is far from realistic. Not only mechanical and electronical inaccuracies inevitable in the construction of antennas, but more importantly the presence of obstacles to the propagation of electromagnetic signals—such as buildings, mountains, or even weather conditions—can lead to areas covered by signal transmission that hardly resemble disks in practice. These considerations motivate that in order to study the described interference reduction problem we abstract from network node positions and circular transmission areas.

In our analysis we formalize the task of reducing interference as a combinatorial optimization problem. For this purpose we model the transmission range of a base station having chosen a specific transmission power level as a set containing exactly all clients covered thereby. The totality of transmission ranges selectable by all base stations is consequently modeled as a collection of client sets. More formally, this yields the *Minimum Membership Set Cover (MMSC)* problem: Given a set of elements U (modeling

clients) and a collection S of subsets of U (transmission ranges), choose a solution $S' \subseteq S$ such that every element occurs in at least one set in S' (maintain network availability) and that the *membership* $M(u, S')$ of any element u with respect to S' is minimal, where $M(u, S')$ is defined as the number of sets in S' in which u occurs (interference).[3]

Having defined this formalization, we show in this paper—by reduction from the related Minimum Set Cover problem—that the MMSC problem is NP-complete and that no polynomial time algorithm exists with approximation ratio less than $\ln n$ unless $NP \subset TIME(n^{\mathrm{O}(\log \log n)})$. In a second part of the paper we present a probabilistic algorithm based on linear programming relaxation and derandomization asymptotically matching this lower bound, particularly yielding an approximation ratio in $\mathrm{O}(\ln n)$.

The paper is organized as follows: Discussing related work in Section 2, we formally define the MMSC problem in Section 3. Section 4 contains a description of the lower bound with respect to approximability of the MMSC problem. In the subsequent section we describe how the MMSC problem can be formulated as a linear program and provide a $\mathrm{O}(\ln n)$-approximation algorithm for the problem. Section 6 concludes the paper.

2 Related Work

Interference issues in cellular networks have been studied since the early 1980s in the context of frequency division multiplexing: The available network frequency spectrum is divided into narrow channels assigned to cells in a way to avoid interference conflicts. In particular two types of conflicts can occur, adjacent cells using the same channel (cochannel interference) and insufficient frequency distance between channels used within the same cell (adjacent channel interference). Maximizing the reuse of channels respecting these conflicts was generally studied by means of the combinatorial problem of conflict graph coloring using a minimum number of colors. The settings in which this problem was considered are numerous and include hexagon graphs, geometric intersection graphs (such as unit disk graphs), and planar graphs, but also (non-geometric) general graphs. In addition both static and dynamic (or on-line) approaches were studied [11]. The fact that channel separation constraints can depend on the distance of cells in the conflict graph was analyzed by means of graph labeling [6]. The problem of frequency assignment was tackled in a different way in [2] exploiting the observation that in every region of an area covered by the communication network it is sufficient that exactly one base station with a unique channel can be heard. As mentioned, all these studied models try to avoid interference conflicts occurring when using frequency division multiplexing. In contrast, the problem described in this paper assumes a different approach in aiming at interference reduction by having the base stations choose suitable transmission power levels.

The problem of reducing interference is formalized in a combinatorial optimization problem named *Minimum Membership Set Cover*. As suggested by its name, at first sight its formulation resembles closely the long-known and well-studied *Minimum Set Cover (MSC)* problem, where the number of sets chosen to cover the given elements

[3] Note that naturally, for each base station, the client set corresponding to a particular power level contains all sets corresponding to lower power levels. Thus, we can assume that w.l.o.g., only one client set is chosen for each base station.

is to be minimized [7]. That the MMSC and the MSC problems are however of different nature can be concluded from the following observation: For any MSC instance consisting of n elements, a greedy algorithm approximates the optimal solution with an approximation ratio at most $H(n) \leq \ln n + 1$ [7], which has later been shown to be tight up to lower order terms unless $NP \subset TIME(n^{O(\log \log n)})$ [3,10]. For the MMSC problem in contrast, there exist instances where the same greedy algorithm fails to achieve *any* nontrivial approximation of the optimal solution.

The approximation algorithm for the MMSC problem introduced in this paper is based on the formulation of a given instance as a linear program. Solving this linear program yields values subsequently interpreted as probabilities with which to randomly decide for every set in S whether it should belong to the solution. This technique, commonly known as randomized rounding was proposed in [12]. Also derandomization based on the method of conditional probabilities—the technique exploited to transform the above probabilistic algorithm into a deterministic one—was introduced in [12] and extended as well as improved in [13].

In the context of network traffic congestion, [9] considered a problem similar to our analysis of the MMSC problem in that linear program relaxation was employed to minimize a maximum value.

3 Minimum Membership Set Cover

As described in the introduction, the problem considered in this paper is to assign to each base station a transmission power level such that interference is minimized while all clients are covered. For our analysis we formalize this problem by introducing a combinatorial optimization problem referred to as *Minimum Membership Set Cover*. In particular, clients are modeled as elements and the transmission range of a base station given a certain power level is represented as the set of thereby covered elements. In the following, we first define the membership of an element given a collection of sets:

Definition 1 (Membership) *Let U be a finite set of elements and S be a collection of subsets of U. Then the membership $M(u, S)$ of an element u is defined as $|\{T \mid u \in T, T \in S\}|$.*

Informally speaking, MMSC is identical to the MSC problem apart from the minimization function. Where MSC minimizes the total number of sets, MMSC tries to minimize element membership. Particularly, MMSC can be defined as follows:

Definition 2 (Minimum Membership Set Cover) *Let U be a finite set of elements with $|U| = n$. Furthermore let $S = \{S_1, \ldots, S_m\}$ be a collection of subsets of U such that $\bigcup_{i=1}^{m} S_i = U$. Then Minimum Membership Set Cover (MMSC) is the problem of covering all elements in U with a subset $S' \subseteq S$ such that $\max_{u \in U} M(u, S')$ is minimal.*[4]

[4] Besides minimizing the *maximal* membership value over all elements, also minimization of the *average* membership value can be considered a reasonable characterization of the interference reduction problem. The fact however that—given a solution S'—the sum of all membership values equals the sum of the cardinalities of the sets in S' shows that this min-average variant is identical to the Weighted Set Cover [1] problem with the set weights corresponding to their cardinalities.

Note that—as motivated in the introduction—the problem statement does not require the collection of subsets S to reflect geometric positions of network nodes. For a given problem instance to be valid, $\bigcup_{i=1}^{m} S_i = U$ is sufficient.

4 Problem Complexity

In this section we address the complexity of the *Minimum Membership Set Cover* problem. We show that MMSC is NP-complete and therefore no polynomial time algorithm exists that solves MMSC unless $P = NP$.

Theorem 1. *MMSC is NP-complete.*

Proof. We will prove that MMSC is NP-complete by reducing MSC to MMSC. Consider an MSC instance (U, S) consisting of a finite set of elements U and a collection S of subsets of U. The objective is to choose a subset S' with minimum cardinality from S such that the union of the chosen subsets of U contains all elements in U.

We now define a set $\widetilde{U}$ by adding a new element e to U, construct a new collection of sets $\widetilde{S}$ by inserting e into all sets in S, and consider $(\widetilde{U}, \widetilde{S})$ as an instance of MMSC. Since element e is in every set in $\widetilde{S}$, it follows that e is an element with maximum membership in the solution S' of MMSC. Moreover, the membership of e in S' is equal to the number of sets in the solution. Therefore MMSC minimizes the number of sets in the solution by minimizing the membership of e. Consequently we obtain the solution for MSC of the instance (U, S) by solving MMSC for the instance $(\widetilde{U}, \widetilde{S})$ and extracting element e from all sets in the solution.

We have shown a reduction from MSC to MMSC, and therefore the latter is NP-hard. Since solutions for the decision problem of MMSC are verifiable in polynomial time, it is in NP, and consequently the MMSC decision problem is also NP-complete. □

Now that we have proved MMSC to be NP-complete and therefore not to be optimally computable within polynomial time unless $P = NP$, the question arises, how closely MMSC can be approximated by a polynomial time algorithm. This is partly answered with the following lower bound.

Theorem 2. *There exists no polynomial time approximation algorithm for MMSC with an approximation ratio less than $(1 - \mathrm{o}(1)) \ln n$ unless $NP \subset TIME(n^{\mathrm{O}(\log \log n)})$.*

Proof. The reduction from MSC to MMSC in the proof of Theorem 1 is approximation-preserving, that is, it implies that any lower bound for MSC also holds for MMSC. In [3] it is shown that $\ln n$ is a lower bound for the approximation ratio of MSC unless $NP \subset TIME(n^{\mathrm{O}(\log \log n)})$. Thus, $\ln n$ is also a lower bound for the approximation ratio of MMSC. □

5 Approximating MMSC by LP Relaxation

In the previous section a lower bound of $\ln n$ for the approximability of the MMSC problem by means of polynomial time approximation algorithms has been established. In this section we show how to obtain a $\mathrm{O}(\log n)$-approximation using LP relaxation techniques.

5.1 LP Formulation of MMSC

We first derive the integer linear program which describes the MMSC problem. Let $S' \subseteq S$ denote a subset of the collection S. To each $S_i \in S$ we assign a variable $x_i \in \{0,1\}$ such that $x_i = 1 \Leftrightarrow S_i \in S'$. For S' to be a set cover, it is required that for each element $u_i \in U$, at least one set S_j with $u_i \in S_j$ is in S'. Therefore, S' is a set cover of U if and only if for all $i = 1, ..., n$ it holds that $\sum_{S_j \ni u_i} x_j \geq 1$. For S' to be minimal in the number of sets that cover a particular element, we need a second set of constraints. Let z be the maximum membership over all elements caused by the sets in S'. Then for all $i = 1, ..., n$ it follows that $\sum_{S_j \ni u_i} x_j \leq z$. The MMSC problem can consequently be formulated as the integer program $\mathrm{IP}_{\mathrm{MMSC}}$:

$$
\begin{aligned}
&\text{minimize } z \\
&\text{subject to} \quad \sum_{S_j \ni u_i} x_j \geq 1 \qquad i = 1, ..., n \\
&\qquad\qquad \sum_{S_j \ni u_i} x_j \leq z \qquad i = 1, ..., n \\
&\qquad\qquad x_j \in \{0,1\} \qquad j = 1, ..., m
\end{aligned}
$$

By relaxing the constraints $x_j \in \{0,1\}$ to $x'_j \geq 0$, we obtain the linear program $\mathrm{LP}_{\mathrm{MMSC}}$. The integer program $\mathrm{IP}_{\mathrm{MMSC}}$ yields the optimal solution z^* for an MMSC problem. The linear program LP_{MMSC} therefore results in a fractional solution z' with $z' \leq z^*$, since we allow the variables x'_j to be in [0,1].

5.2 Randomized Rounding

In [12] and [13], randomized rounding was introduced for covering and packing problems. In the following, we show that this technique can also be applied to solve $\mathrm{IP}_{\mathrm{MMSC}}$ resulting in an almost optimal algorithm. We present an efficient $(1+\mathrm{O}(1/\sqrt{z'}))(\ln(n)+1)$-approximation algorithm for the MMSC problem. Given an MMSC instance (U, S), we first solve the linear program $\mathrm{LP}_{\mathrm{MMSC}}$ corresponding to (U, S), yielding a vector $\underline{x}'$ and z' and then apply randomized rounding in order to obtain an integer solution. Consider the following "simple" randomized rounding scheme. We compute an integer solution $\underline{x} \in \{0,1\}^m$ by setting

$$
x_i := \begin{cases} 1 & \text{with probability } p_i := \min\{1, \alpha x'_i\} \text{ for a value } \alpha \geq 1 \\ 0 & \text{otherwise} \end{cases}
$$

independently for each $i \in \{1, \ldots, m\}$. Let $\mathcal{A}_i$ be the *bad* event that the i^{th} element is *not* covered.

Lemma 1. *The probability that the i^{th} element remains uncovered is*

$$
\mathbb{P}(\mathcal{A}_i) \;=\; \prod_{S_j \ni u_i} (1 - p_j) \;< e^{-\alpha}.
$$

Proof. Let m_i be the number of sets containing element i ($m_i = |\{j|S_j \ni u_i\}|$). By the "means inequality", we have

$$\mathbb{P}(\mathcal{A}_i) = \prod_{S_j \ni u_i} (1 - p_j) \leq \left(1 - \frac{\sum_{S_j \ni u_i} p_j}{m_i}\right)^{m_i} \leq \left(1 - \frac{\alpha}{m_i}\right)^{m_i} < e^{-\alpha}.$$

Note that $\sum_{S_j \ni u_i} p_j \geq \alpha$ only holds if all $p_j < 1$. We can safely make this assumption because $p_j = 1$ for some $S_j \ni u_i$ makes $\mathbb{P}(\mathcal{A}_i) = 0$. □

Let $\mathcal{B}_i$ be the *bad* event that the i^{th} element is covered by *more* than $\alpha\beta z'$ sets for some $\beta \geq 1$.

Lemma 2. *The probability that the i^{th} element is covered more than $\alpha\beta z'$ times is*

$$\mathbb{P}(\mathcal{B}_i) < \frac{1}{\beta^{\alpha\beta z'}} \cdot \prod_{S_j \ni u_i} [1 + (\beta - 1)p_j] \leq \left(\frac{e^{\beta-1}}{\beta^\beta}\right)^{\alpha z'}.$$

Proof. We use a Chernoff-type argument. For $t = \ln \beta > 0$, we have

$$\begin{aligned}
\mathbb{P}(\mathcal{B}_i) = \mathbb{P}\left(\sum_{S_j \ni u_i} x_j > \alpha\beta z'\right) &= \mathbb{P}\left(e^{t \cdot \sum_{S_j \ni u_i} x_j} > e^{t\alpha\beta z'}\right) \\
&< \frac{\mathrm{E}\left[e^{t \cdot \sum_{S_j \ni u_i} x_j}\right]}{e^{t\alpha\beta z'}} = \frac{1}{e^{t\alpha\beta z'}} \cdot \prod_{S_j \ni u_i} \left[p_j e^t + 1 - p_j\right] \\
&= \frac{1}{\beta^{\alpha\beta z'}} \cdot \prod_{S_j \ni u_i} [1 + (\beta - 1)p_j] \leq \frac{1}{\beta^{\alpha\beta z'}} \cdot \prod_{S_j \ni u_i} e^{(\beta-1)p_j} \leq \left(\frac{e^{\beta-1}}{\beta^\beta}\right)^{\alpha z'}.
\end{aligned}$$

The inequality in the first line results by application of the Markov inequality. The equations in the second line hold because of the independence of the x_i and because $t = \ln \beta$, respectively. For the inequalities in the last line, we apply $1 + x \leq e^x$ and $\sum_{S_j \ni u_i} p_j \leq \alpha z'$. □

In the following, we denote the probability upper bounds given by Lemmas 1 and 2 by A_i and B_i:

$$A_i := \prod_{S_j \ni u_i} (1 - p_j) \quad \text{and} \quad B_i := \frac{1}{\beta^{\alpha\beta z'}} \cdot \prod_{S_j \ni u_i} [1 + (\beta - 1)p_j].$$

In order to bound the probability for any bad event to occur, we define a function P as follows:

$$P(p_1, \ldots, p_m) := 2 - \prod_{i=1}^{n} (1 - A_i) - \prod_{i=1}^{n} (1 - B_i).$$

Lemma 3. *The probability that any element is not covered or covered more than $\alpha\beta z'$ times is upper-bounded by $P(p_1, \ldots, p_m)$:*

$$\mathbb{P}\left(\bigcup_{i=1}^{n} \mathcal{A}_i \cup \bigcup_{i=1}^{n} \mathcal{B}_i\right) < P(p_1, \ldots, p_m).$$

Proof. It is sufficient to prove that

$$\mathbb{P}\left(\bigcup_{i=1}^{n} \mathcal{A}_i\right) \leq 1 - \prod(1 - \mathbb{P}(\mathcal{A}_i)) \quad \text{and} \quad \mathbb{P}\left(\bigcup_{i=1}^{n} \mathcal{B}_i\right) \leq 1 - \prod(1 - \mathbb{P}(\mathcal{B}_i)). \quad (1)$$

The lemma then follows by Lemmas 1 and 2. If the events $\mathcal{A}_i$ and $\mathcal{B}_i$ were independent, the first and second inequality of (1) would hold with equality, respectively. Hence, we have to show that the dependence of the events can only help us. As shown in [13], the, the complementary events $\overline{\mathcal{A}}_i$ are positively correlated, that is, the probability of $\overline{\mathcal{A}}_i$ ($\mathcal{A}_i$ does not occur) increases under the condition that any subset of $\{\overline{\mathcal{A}}_1, \ldots, \overline{\mathcal{A}}_n\}$ occurs. This positive correlation follows from Harris-Kleitman inequality [5,8], which is a special case of the FKG inequality [4]. Hence, the first inequality of (1) follows. For the events $\mathcal{B}_i$ exactly the same argumentation holds. □

In the following we show that if α and β are chosen appropriately, $P(p_1, \ldots, p_m)$ is always less than 1.

Lemma 4. *When setting $\alpha = \ln(n) + 1$, then for $\beta = 1 + \max\{\sqrt{3/z'}, 3/z'\}$, we have $P(p_1, \ldots, p_m) < 4/5$.*

Proof. By Lemmas 1 and 2, we have

$$P(p_1, \ldots, p_m) < 2 - \left(1 - e^{-\alpha}\right)^n - \left(1 - \left(\frac{e^{\beta-1}}{\beta^\beta}\right)^{\alpha z'}\right)^n.$$

In order to have $P < 4/5$, it therefore suffices to choose α and β such that

$$\left(1 - e^{-\alpha}\right)^n \geq \frac{3}{5} \quad \text{and} \quad \left(1 - \left(\frac{e^{\beta-1}}{\beta^\beta}\right)^{\alpha z'}\right)^n \geq \frac{3}{5}. \quad (2)$$

For $\alpha \geq \ln n + 1$, we get $(1 - e^{-\alpha})^n \geq 3/5$ and therefore the first inequality of (2) is fulfilled. The second inequality of (2) can be transformed into a simpler form using the following inequalities:

$$\frac{e^{\beta-1}}{\beta^\beta} \leq \begin{cases} e^{-(\beta-1)^2/3} & \text{for } 1 \leq \beta \leq 2, \\ e^{-(\beta-1)/3} & \text{for } \beta > 2. \end{cases} \quad (3)$$

If we choose $\beta = 1 + \sqrt{3/z'}$, for $z' \geq 3$, we have $\beta \leq 2$ and therefore by (3), the second inequality of (2) simplifies to

$$\left(1 - \left(\frac{e^{\beta-1}}{\beta^\beta}\right)^{\alpha z'}\right)^n \geq \left(1 - e^{-\alpha z'(\beta-1)^2/3}\right)^n = \left(1 - e^{-\alpha}\right)^n.$$

For $z' < 3$, we can set $\beta = 1 + 3/z' \geq 2$ and proceed analogously using the second case of (3). □

Lemmas 1–4 lead to the following randomized algorithm for the MMSC problem. As a first step, the linear program $\mathrm{LP}_{\mathrm{MMSC}}$ has to be solved. Then, all x'_i are rounded to integer values $x_i \in \{0, 1\}$ using the described randomized rounding scheme with $\alpha = \ln n + 1$. The rounding is repeated until the solution is feasible (all elements are covered) and the membership of the integer solution deviates from the fractional membership z' by at most a factor $\alpha\beta$ for $\beta = 1 + \max\{3/z', \sqrt{3/z'}\}$. Each time, the probability to be successful is at least $1/5$ and therefore, the probability of not being successful decreases exponentially in the number of trials.

5.3 Derandomization

We will now show that $P(p_1, \ldots, p_m)$ is a pessimistic estimator [12,13] and that therefore, the algorithm described at the end of the previous section can be derandomized. That is, P is an upper bound on the probability of obtaining a "bad" solution, $P < 1$ (P is a probabilistic proof that a "good" solution exists), and the p_i can be set to 0 or 1 without increasing P. The first two properties follow by Lemmas 3 and 4, the third property is shown by the following lemma.

Lemma 5. *For all i, either setting p_i to 0 or setting p_i to 1 does not increase P:*

$$P(p_1, \ldots, p_m) \;\geq\; \min\{P(\ldots, p_{i-1}, 0, p_{i+1}, \ldots), P(\ldots, p_{i-1}, 1, p_{i+1}, \ldots)\}.$$

Proof. We prove the lemma by showing that P is a concave function of p_i:

$$P(p_1, \ldots, p_m) \geq (1-p_i)P(\ldots, p_{i-1}, 0, p_{i+1}, \ldots) + p_i P(\ldots, p_{i-1}, 1, p_{i+1}, \ldots). \quad (4)$$

If all probabilities except p_i are fixed, A_j and B_j are functions of p_i. We define

$$\overline{A}_k(p_i) \;:=\; \prod_{j=1}^{k}(1 - A_j) \quad \text{and} \quad \overline{B}_k(p_i) \;:=\; \prod_{j=1}^{k}(1 - B_j).$$

In order to obtain (4), we prove that

$$\overline{A}_k(p_i) \leq (1 - p_i)\overline{A}_k(0) + p_i\overline{A}_k(1) \quad \text{and} \quad \overline{B}_k(p_i) \leq (1 - p_i)\overline{B}_k(0) + p_i\overline{B}_k(1) \quad (5)$$

for all $k \in [0, n]$ by induction over k. For $k = 0$, we have $\overline{A}_0(p_i) = \overline{B}_0(p_i) = 1$ and therefore (5) holds. The induction step from k to $k + 1$ depends on whether element $k + 1$ is in set S_i. If element $k + 1$ is not in set S_i, A_{k+1} and B_{k+1} do not depend on p_i and (5) follows from the induction hypothesis. It remains to prove the interesting case where element $k + 1$ is contained in set S_i. We first consider the inequality for $\overline{A}_{k+1}(p_i)$. When p_i is set to 1, A_{k+1} becomes 0. If p_i is set to 0, the factor $1 - p_i$ in A_{k+1} is replaced by 1 and therefore A_{k+1} becomes

$$A_{k+1,p_i=0} = \prod_{S_j \in u_{k+1} \setminus S_i} (1 - p_j) \;=\; \frac{A_{k+1}}{1 - p_i}.$$

We therefore have

$$\begin{aligned}(1-p_i)\overline{A}_{k+1}(0) + p_i\overline{A}_{k+1}(1) &= (1-p_i)\overline{A}_k(0)\cdot\left(1-\frac{A_{k+1}}{1-p_i}\right) + p_i\overline{A}_k(1)\cdot 1\\ &= (1-p_i)\overline{A}_k(0) + p_i\overline{A}_k(1) - \overline{A}_k(0)\cdot A_{k+1}\\ &\geq \overline{A}_k(p_i)(1-A_{k+1}) \;=\; \overline{A}_{k+1}(p_i).\end{aligned}$$

The inequality in the third line follows from the induction hypothesis and from $\overline{A}_k(0) \leq \overline{A}_k(p_i)$. For $\overline{B}_{k+1}(p_i)$, setting p_i to 0 and 1 replaces the factor $1+(\beta-1)p_i$ in B_{k+1} by 1 and β, respectively:

$$B_{k+1,p_i=0} = \frac{B_{k+1}}{1+(\beta-1)p_i} \quad \text{and} \quad B_{k+1,p_i=1} = \frac{\beta B_{k+1}}{1+(\beta-1)p_i}.$$

Thus, we get

$$\begin{aligned}&(1-p_i)\overline{B}_{k+1}(0) + p_i\overline{B}_{k+1}(1)\\ &= (1-p_i)\overline{B}_k(0)\cdot\left(1-\frac{B_{k+1}}{1+(\beta-1)p_i}\right) + p_i\overline{B}_k(1)\cdot\left(1-\frac{\beta B_{k+1}}{1+(\beta-1)p_i}\right)\\ &= \left((1-p_i)\overline{B}_k(0) + p_i\overline{B}_k(1)\right)\cdot\left(1-\frac{B_{k+1}}{1+(\beta-1)p_i}\right) - \frac{B_{k+1}p_i(\beta-1)\overline{B}_k(1)}{1+(\beta-1)p_i}\\ &\geq \overline{B}_k(p_i)\cdot\left(1-\frac{B_{k+1}}{1+(\beta-1)p_i}\right) - \frac{B_{k+1}p_i(\beta-1)\overline{B}_k(p_i)}{1+(\beta-1)p_i}\\ &= \overline{B}_k(p_i)(1-B_{k+1}) \;=\; \overline{B}_{k+1}(p_i).\end{aligned}$$

The inequality in the fourth line follows from the induction hypothesis and from $\overline{B}_k(1) \leq \overline{B}_k(p_i)$. □

Lemmas 3, 4, and 5 lead to an efficient deterministic approximation algorithm for the MMSC problem. First, the linear program $\mathrm{LP}_{\mathrm{MMSC}}$ has to be solved. The probabilites p_i are determined as described in the last section. For α and β as in Lemma 4, $P(p_1,\ldots,p_m) < 4/5$. The probabilities p_i are now set to 0 or 1 such that $P(p_1,\ldots,p_m)$ remains smaller than $4/5$. This is possible by Lemma 5. When all $p_i \in \{0,1\}$, we have an integer solution for $\mathrm{IP}_{\mathrm{MMSC}}$. The probability that not all elements are covered or that the membership is larger than $\alpha\beta z'$ is smaller than $P < 4/5$. Because all p_i are 0 or 1, this probability must be 0. Hence, the computed $\mathrm{IP}_{\mathrm{MMSC}}$-solution is an $\alpha\beta$-approximation for MMSC:

Theorem 3. *For any MMSC instance, there exists a deterministic polynomial-time approximation algorithm with an approximation ratio of* $(1 + \mathrm{O}(1/\sqrt{z'}))(\ln(n)+1)$.

6 Conclusion

Interference reduction in cellular networks is studied in this paper by means of formalization with the *Minimum Membership Set Cover* problem. Although this combinatorial optimization problem appears to be a natural and simply describable problem

in the context of set covering, it has—to the best of our knowledge—not been studied before. We show using approximation-preserving reduction from the Minimum Set Cover problem that MMSC is not only NP-hard, but also that no polynomial-time algorithm can approximate the optimal solution more closely than up to a factor $\ln n$ unless $NP \subset TIME(n^{\mathrm{O}(\log \log n)})$. In a second part of the paper this lower bound is shown to be asymptotically matched by an algorithm making use of linear programming relaxation techniques.

Finally, the question remains as an open problem, whether there exists a simpler greedy algorithm—considering interference increase during its execution—with the same approximation quality.

References

1. V. Chvátal. A Greedy Heuristic for the Set-Covering Problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
2. G. Even, Z. Lotker, D. Ron, and S. Smorodinsky. Conflict-Free Colorings of Simple Geometric Regions with Applications to Frequency Assignment in Cellular Networks. In 43^{rd} *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2002.
3. U. Feige. A Threshold of ln n for Approximating Set Cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.
4. C. Fortuin, P. Kasteleyn, and J. Ginibre. Correlations Inequalities on Some Partially Ordered Sets. *Comm. Math. Phys.*, 22:89–103, 1971.
5. T. E. Harris. A Lower Bound for the Critical Probability in a Certain Percolation Process. *Proc. Cambridge Philos. Soc.*, 60:13–20, 1960.
6. J. Janssen. Channel Assignment and Graph Labeling. In I. Stojmenovic, editor, *Handbook of Wireless Networks and Mobile Computing*, chapter 5, pages 95–117. John Wiley & Sons, Inc., 2002.
7. D. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.
8. D. Kleitman. Families of Non-Disjoint Subsets. *Journal of Combinatorial Theory*, 1:153–155, 1966.
9. C.-J. Lu. A Deterministic Approximation Algorithm for a Minimax Integer Programming Problem. In *Proc. of the 10th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 663–668, 1999.
10. C. Lund and M. Yannakakis. On the Hardness of Approximating Minimization Problems. *Journal of the ACM*, 41(5):960–981, 1994.
11. L. Narayanan. Channel Assignment and Graph Multicoloring. In I. Stojmenovic, editor, *Handbook of Wireless Networks and Mobile Computing*, chapter 4, pages 71–94. John Wiley & Sons, Inc., 2002.
12. P. Raghavan and C. Thompson. Randomized Rounding: A Technique for Provably Good Algorithms and Algorithmic Proofs. *Combinatorica*, 7(4):365–374, 1987.
13. A. Srinivasan. Improved Approximations of Packing and Covering Problems. In *Proc. of the 27th ACM Symposium on Theory of Computing*, pages 268–276, 1995.