Research Collection

Report

On coverage in infrastructure networks

Author(s):
Kuhn, Fabian; Rickenbach, Pascal von; Wattenhofer, Roger; Welzl, Emo; Zollinger, Aaron

ETH Library

On Coverage in Infrastructure Networks

Fabian Kuhn, Pascal von Rickenbach, Roger Wattenhofer, Emo Welzl, Aaron Zollinger
{kuhn, rickenbach, wattenhofer, welzl, zollinger}@inf.ethz.ch
Department of Computer Science, ETH Zurich, 8092 Zurich, Switzerland

Abstract

For mobile distributed tasks, the infrastructure 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 not only asymptotically matches the lower bound with high probability but also performs well in practical networks.

Keywords: Interference, cellular networks, combinatorial optimization, approximation algorithms, linear programming relaxation.

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. Cellu-lar networks can however more broadly be considered a type of infrastructure for 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 over-lapping 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, 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
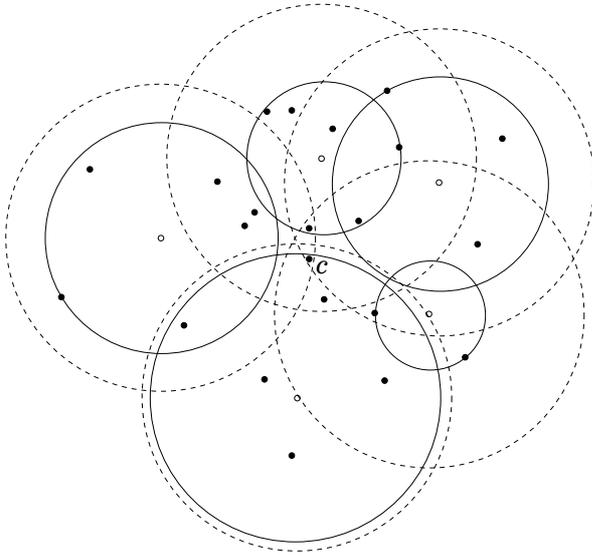
Figure 1: If the base stations (hollow points) are assigned identical transmission power levels (dashed circles), client c experiences high interference, since it is covered by all base stations. Interference can be reduced by assigning appropriate power values (solid circles), such that all clients are covered by at most two base stations.

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 our analysis we formalize this task 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* $M(e, S')$ of any element e with respect to S' is minimal, where $M(e, S')$ is defined as the number of sets in S' in which e occurs (interference).

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 $NP \subset TIME(n^{O(\log \log n)})$. We additionally present a probabilistic algorithm based on linear programming relaxation asymptotically matching this lower bound, particularly yielding an approximation ratio in $O(\log n)$ with high probability. Furthermore we study how the presented algorithm performs on practical network instances.

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 $O(\log n)$-approximation algorithm for the problem. The behavior of the proposed algorithm in practical networks is the subject of Section 6. Section 7 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 net-

work 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 [9]. The fact that channel separation constraints can depend on the distance of cells in the conflict graph was studied by means of graph labeling [4]. The problem of frequency assignment is 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 is to be minimized [5]. 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 $H(n) \leq \ln n + 1$ [5], which has later been shown to be tight up to lower order terms unless $NP \subset TIME(n^{O(\log \log n)})$ [3, 7]. For the MMSC problem in contrast, there exist instances where the same greedy algorithm fails to achieve *any* nontrivial approximation of the optimal solution.

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 $M(e, S)$ of an element e is defined as $|\{T \mid e \in T, T \in S\}|$.*

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 $\max_{e \in U} M(e, S')$ is minimal.*

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 $\ln n$ unless $NP \subset TIME(n^{O(\log \log n)})$.*

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 $NP \subset TIME(n^{O(\log \log n)})$. 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 $O(\log n)$-approximation with high probability[1] using LP relaxation techniques. For an introduction to linear programming see for instance [1].

5.1 LP Formulation of MMSC

We first derive the integer linear program which describes the MMSC problem and then formulate the linear program that relaxes the integrality constraints.

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 $\sum_{S_j : u_i \in S_j} x_j \geq 1$. 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

---

[1]Throughout the paper, an event E occurring "with high probability" stands for $Pr[E] = 1 - O\left(\frac{1}{n}\right)$.

all elements caused by the sets in S'. Then for all $i = 1, ..., n$ it follows that $\sum_{S_j : u_i \in S_j} x_j \leq z$. The MMSC problem can consequently be formulated as the integer program IP_{MMSC}:

minimize z

subject to
$$\sum_{S_j : u_i \in S_j} x_j \geq 1 \qquad i = 1, ..., n$$
$$\sum_{S_j : u_i \in S_j} x_j \leq z \qquad i = 1, ..., n$$
$$x_j \in \{0, 1\} \qquad j = 1, ..., m$$

By relaxing the constraints $x_j \in \{0, 1\}$ to $x'_j \geq 0$, we obtain the following linear program LP_{MMSC}:

minimize z

subject to
$$\sum_{S_j : u_i \in S_j} x'_j \geq 1 \qquad i = 1, ..., n$$
$$\sum_{S_j : u_i \in S_j} x'_j \leq z \qquad i = 1, ..., n$$
$$x'_j \geq 0 \qquad j = 1, ..., m$$

The integer program IP_{MMSC} yields the optimal solution z^* for an MMSC problem. The derived linear program LP_{MMSC} therefore obtains a fractional solution z' with $z' \leq z^*$, since we allow the variables x'_j to be in [0,1].

5.2 Algorithm and Analysis

We will now present a $\text{O}(\log n)$-approximation algorithm, referred to as $\mathcal{A}_{\text{MMSC}}$, for the MMSC problem. Given an MMSC instance (U, S), the algorithm first solves the linear program LP_{MMSC} corresponding to (U, S). In a second step, $\mathcal{A}_{\text{MMSC}}$ performs randomized rounding (see [10]) on a feasible solution vector $\underline{x}'$ for LP_{MMSC}, in order to derive a vector $\underline{x}$ with $x_i \in \{0, 1\}$. Finally it is ensured that $\underline{x}$ is a feasible solution for IP_{MMSC} and consequently a set cover.

For the analysis of $\mathcal{A}_{\text{MMSC}}$ the following two mathematical facts are required. Their proofs are

---

Algorithm $\mathcal{A}_{\text{MMSC}}$

Input: an MMSC instance (U, S)
1: compute solution vector $\underline{x}'$ to the linear program LP_{MMSC} corresponding to (U, S)
2: $p_i := \min\{1, \ x'_i \cdot \log n\}$
3: $x_i := \begin{cases} 1 & \text{with probability } p_i \\ 0 & \text{otherwise} \end{cases}$
4: **for all** $u_i \in U$ **do**
5:    **if** $\sum_{S_j : u_i \in S_j} x_j = 0$ **then**
6:       set $x_j = 1$ for any j such that $u_i \in S_j$
7:    **end if**
8: **end for**
Output: MMSC solution S' corresponding to $\underline{x}$

---

omitted and can be found in mathematical text books.

Fact 1. (Means Inequality) *Let $\mathcal{A} \subset \mathbb{R}^+$ be a set of positive real numbers. The product of the values in $\mathcal{A}$ can be upper-bounded by replacing each factor with the arithmetic mean of the elements of $\mathcal{A}$:*

$$\prod_{x \in \mathcal{A}} x \ \leq \ \left(\frac{\sum_{x \in \mathcal{A}} x}{|\mathcal{A}|} \right)^{|\mathcal{A}|}.$$

Fact 2. *For all n, t, such that $n \geq 1$ and $|t| \leq n$,*

$$e^t \left(1 - \frac{t^2}{n} \right) \leq \left(1 + \frac{t}{n} \right)^n \ \leq \ e^t.$$

We prove $\mathcal{A}_{\text{MMSC}}$ to be a $\text{O}(\log n)$-approximation algorithm for IP_{MMSC} in several steps. We first show that the membership of an element in U after the randomized rounding step of $\mathcal{A}_{\text{MMSC}}$ is bounded with high probability.

Lemma 3. *The membership of an element u_i after Line 3 of $\mathcal{A}_{\text{MMSC}}$ is at most $2e \log n \cdot z^*$ with high probability.*

Proof. The optimal solution of LP_{MMSC} leads to fractional values x'_j and does not admit a straight-forward choice of the sets S_j. Using randomized rounding, $\mathcal{A}_{\text{MMSC}}$ converts the fractional solution

to an integral solution S'. In Line 3, a set S_j is chosen to be in S' with probability $x'_j \cdot \log n$. Thus, the expected membership of an element u_i is

$$
\begin{aligned}
E[M(u_i, S')] &= \sum_{S_j : u_i \in S_j} x'_j \cdot \log n \\
&\leq \log n \cdot z'. \quad (1)
\end{aligned}
$$

The last inequality follows directly from the second set of constraints of LP_{MMSC}. Since $z' \leq z^*$, it follows that the expected membership for u_i is at most $\log n \cdot z^*$. Now we need to ensure that, with high probability, u_i is not covered too often. Since randomized rounding can be modeled as Poisson trials, we are able to use a Chernoff bound [8]. Let Y_i be a random variable denoting the membership of u_i with expected value $\mu = E[M(u_i, S')]$. Applying the Chernoff bound we derive

$$
Pr\left[Y_i \geq (1 + \delta)\mu\right] < \left(\frac{e^\delta}{(1 + \delta)^{(1+\delta)}}\right)^\mu.
$$

Choosing $\delta \geq 2e - 1$, the right hand side of the inequality simplifies to

$$
\begin{aligned}
\left(\frac{e^\delta}{(1 + \delta)^{(1+\delta)}}\right)^\mu &\leq \left(\frac{e^\delta}{(2e)^{(1+\delta)}}\right)^\mu \\
&< \left(\frac{e^\delta}{(2e)^\delta}\right)^\mu \\
&= 2^{-\delta\mu}. \quad (2)
\end{aligned}
$$

Since the above Chernoff bound corresponds to the upper tail of the probability distribution of Y_i and as μ is at most $\log n \cdot z^*$, it follows that

$$
Pr\left[Y_i \geq (1 + \delta)\log n \cdot z^*\right] \leq Pr\left[Y_i \geq (1 + \delta)\mu\right].
$$

But for this inequality to hold, only $(1 + \delta)\mu \leq c \log n \cdot z^*$ for some constant c is required. Thus, by setting $(1 + \delta)\mu = c \log n \cdot z^*$ and using Inequality (1), we obtain

$$
\delta\mu \geq (c - 1)\log n \cdot z^*. \quad (3)
$$

Using Inequalities (2) and (3) we can then bound the probability that the membership of u_i is greater than $c \log n \cdot z^*$ as follows:

$$
\begin{aligned}
Pr\left[Y_i \geq c \log n \cdot z^*\right] &< 2^{-\delta\mu} \\
&\leq 2^{-(c-1)\log n \cdot z^*} \\
&= \frac{1}{n^{(c-1)z^*}}. \quad (4)
\end{aligned}
$$

In order to compute c, we again consider the equation $(1 + \delta)\mu = c \log n \cdot z^*$. Solving for δ, we derive

$$
\delta = \frac{c \log n \cdot z^*}{\mu} - 1.
$$

As a requirement for Inequality (2) we demand δ to be greater or equal to $2e - 1$. Furthermore, the right hand side of the inequality is minimal if μ is maximal. Thus, using Inequality (1) we obtain

$$
\frac{c \log n \cdot z^*}{\log n \cdot z^*} - 1 \geq 2e - 1
$$

or $c \geq 2e$. Taking everything together and using $z^* \geq 1$ it follows that

$$
Pr\left[Y_i \geq 2e \log n \cdot z^*\right] < \frac{1}{n^{(2e-1)z^*}} \in \mathrm{O}\left(\frac{1}{n^4}\right).
$$

$\square$

Now we are ready to show that after randomized rounding all elements have membership at most $2e \log n \cdot z^*$ with high probability.

Lemma 4. *The membership of all elements in U after Line 3 of $\mathcal{A}_{\text{MMSC}}$ is at most $2e \log n \cdot z^*$ with high probability.*

Proof. Let E_i be the event that the membership of element u_i after Line 3 of $\mathcal{A}_{\text{MMSC}}$ is greater than $2e \log n \cdot z^*$. Then, the probability that the membership for all elements in U is less than $2e \log n \cdot z^*$ equals

$$
Pr[\bigwedge_{i=1}^{n} \overline{E_i}].
$$

We know from Lemma 3 that the probability $Pr[E_i]$ is less than $1/n^{(2e-1)z^*}$. Since the events

6

are clearly not independent, we cannot apply the product rule. However, it was shown in [11] that

$$Pr[\bigwedge_{i=1}^{n} \overline{E_i}] \geq \prod_{i=1}^{n} Pr[\overline{E_i}]. \qquad (5)$$

We can make use of this bound, since IP_{MMSC} features the *positive correlation* property assumed in [11]. Consequently, setting $\alpha = (2e-1)z^*$ and using Inequality (5), it follows that

$$
\begin{aligned}
Pr[\bigwedge_{i=1}^{n} \overline{E_i}] &\geq \left(1 - \frac{1}{n^{\alpha}}\right)^{n} \\
&\geq \left(1 - \frac{1}{n^{\alpha}}\right)^{\frac{n^{\alpha}-1}{n^{\alpha-1}-\frac{1}{n}}} \\
&\geq e^{-\frac{1}{n^{\alpha-1}-\frac{1}{n}}} \qquad (6) \\
&> 1 - \frac{1}{n^{\alpha-1}-\frac{1}{n}}. \qquad (7)
\end{aligned}
$$

For Inequality (6) we use Fact 2 with $t = -1$, which leads to the inequality

$$e^{-1} \leq (1 - 1/n)^{n-1}.$$

Inequality (7) is derived through Taylor series expansion of (6). Consequently, using $\alpha = (2e-1)z^*$ and $z^* \geq 1$ we obtain

$$Pr[\bigwedge_{i=1}^{n} \overline{E_i}] = 1 - \text{O}\left(\frac{1}{n^3}\right).$$

$\square$

Since $\mathcal{A}_{\text{MMSC}}$ uses randomized rounding, we do not always derive a feasible solution for IP_{MMSC} after Line 3 of the algorithm. That is, there exist elements in U that are not covered by a set in S'. But we can show in the following lemma that each single element is covered with high probability.

Lemma 5. *After Line 3 of $\mathcal{A}_{\text{MMSC}}$, an element u_i in U is covered with high probability.*

Proof. For convenience we define C_i to be the set $\{S_j \mid u_i \in S_j\}$. From LP_{MMSC} we know that $\sum_{S_j \in C_i} x'_j \geq 1$. Thus, it follows that

$$\sum_{S_j \in C_i} p_j \geq \log n. \qquad (8)$$

Let q_i be the probability that an element u_i is contained in none of the sets in S' obtained by randomized rounding, that is, $q_i = Pr[M(u_i, S') = 0]$. Consequently, we have

$$
\begin{aligned}
q_i &= \prod_{S_j \in C_i} (1 - p_j) \\
&\leq \left(1 - \frac{\sum_{S_j \in C_i} p_j}{|C_i|}\right)^{|C_i|} \\
&\leq e^{-\sum_{S_j \in C_i} p_j} \\
&\leq e^{-\log n} \\
&= \frac{1}{n}.
\end{aligned}
$$

The first inequality follows from Fact 1, the second inequality follows from Fact 2, and the third step is derived from Inequality (8).

$\square$

In Lines 4 to 8 of $\mathcal{A}_{\text{MMSC}}$ it is ensured that the final solution S' is a set cover. This is achieved by consecutively including sets in S', until all elements are covered. In the following we show that the additional maximum membership increase caused thereby is bounded with high probability.

Lemma 6. *In Lines 4 to 8 of $\mathcal{A}_{\text{MMSC}}$, the maximum membership in U is increased by at most $\text{O}(\log n)$ with high probability.*

Proof. In order to bound the number of sets added in the considered part of the algorithm, again a Chernoff bound is employed. Let Z be a random variable denoting the number of uncovered elements after Line 3 of $\mathcal{A}_{\text{MMSC}}$. From Lemma 5 we know that an element is uncovered after randomized rounding with probability less than $1/n$.

Then, the expected value μ for Z is less than 1. Using a similar analysis as in Lemma 3, we obtain

$$Pr\left[Z \geq c\right] < 2^{-c+1},$$

where $c \geq 2e$ is required. Setting $c = \log n + 2e$, it follows that

$$Pr\left[Z \geq \log n + 2e\right] \; < \; \frac{2}{n \cdot 4^e} \; \in \; \mathrm{O}\left(\frac{1}{n}\right).$$

The proof is concluded by the observation that each additional set added in the second step of $\mathcal{A}_{\mathrm{MMSC}}$ increases the maximum membership in U by at most one. Since only $\mathrm{O}(\log n)$ elements have to be covered with high probability and as it is sufficient to add one set per element, the lemma follows.[2] □

Now we are ready to prove that $\mathcal{A}_{\mathrm{MMSC}}$ yields a $\mathrm{O}(\log n)$-approximation for $\mathrm{IP}_{\mathrm{MMSC}}$ and consequently also for MMSC.

Theorem 7. *Given an MMSC instance consisting of m sets and n elements, $\mathcal{A}_{\mathrm{MMSC}}$ computes a $\mathrm{O}(\log n)$-approximation with high probability. The running time of $\mathcal{A}_{\mathrm{MMSC}}$ is polynomial in $m \cdot n$.*

Proof. The approximation factor in the theorem directly follows from Lemmas 4 and 6. The running time result is a consequence to the existence of algorithms solving linear programs in time polynomial in the program size [6] and to the fact that $\mathrm{LP}_{\mathrm{MMSC}}$ can be described using -1, 0, and 1 as coefficients only. □

5.3 Alternative Algorithm

In an alternative version of the algorithm, the values $\underline{x}'$ obtained by solving $\mathrm{LP}_{\mathrm{MMSC}}$ can be directly employed as probabilities for randomized rounding (without the additional factor of $\log n$).

In this case randomized rounding is repeated for all sets containing elements not yet covered until resulting in a set cover. With similar arguments as for $\mathcal{A}_{\mathrm{MMSC}}$, it can be shown that this modified algorithm achieves the same approximation factor and that it terminates after repeating randomized rounding at most $\log n$ times, both with high probability.

6 Practical Networks

Whereas the previous section showed that $\mathcal{A}_{\mathrm{MMSC}}$ approximates the optimal solution up to a factor in $\mathrm{O}(\log n)$, this section discusses practical networks. In particular, the algorithms $\mathcal{A}_{\mathrm{MMSC}}$ and $\widetilde{\mathcal{A}}_{\mathrm{MMSC}}$—the alternative algorithm described in Section 5.3—are considered. Since the approximation performance of algorithms is studied, we denote by the *membership of a solution* the minimization function value—that is the maximum membership over all clients—of the corresponding MMSC solution.

The studied algorithms were executed on instances generated by placing base stations and clients randomly according to a uniform distribution on a square field with side length 5 units. Adaptable transmission power values were modeled by attributing to each base station circles with radii 0.25, 0.5, 0.75, and 1 unit; each such circle then contributes one set containing all covered clients to the problem instance thereafter presented to the algorithms.

As shown in the previous section, the approximation factor of the algorithms depends on the number of clients. For this reason the simulations were carried out over a range of client densities. Since the membership value obtained by solving $\mathrm{LP}_{\mathrm{MMSC}}$ lies below the optimal solution and therefore the gap between the algorithm result and the solution of the linear program is an upper bound for the obtained approximation ratio, the $\mathrm{LP}_{\mathrm{MMSC}}$ result z' is also considered.

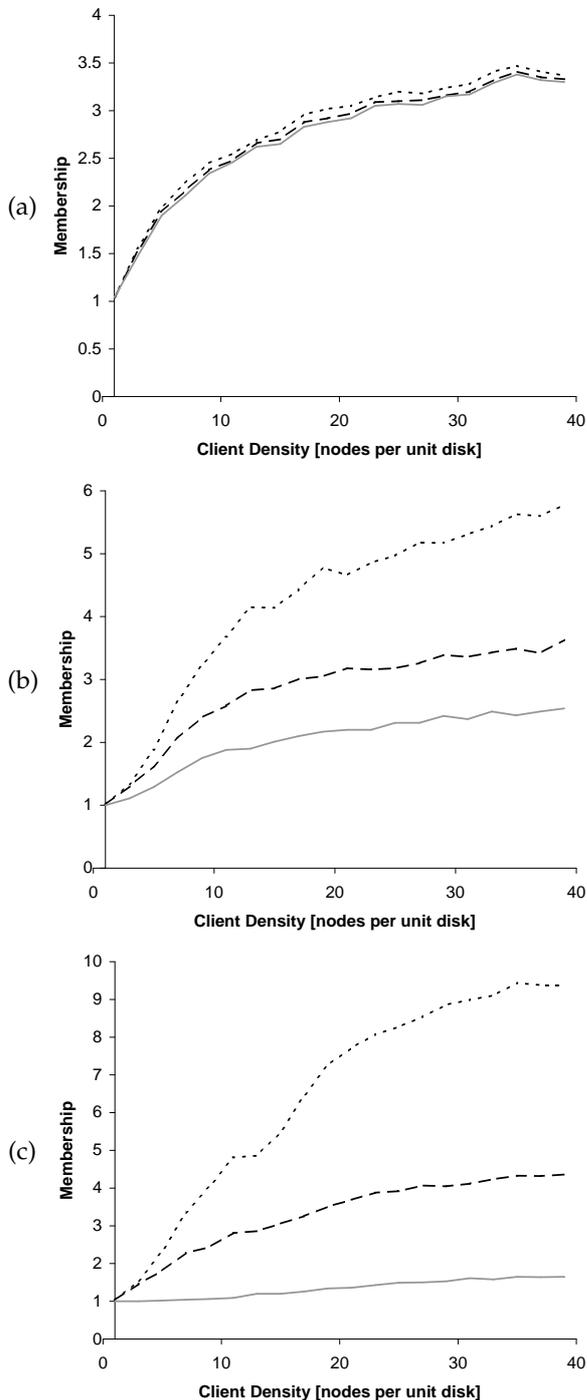For a base-station density of 2 base stations per

(a)

(b)

(c)

Figure 2: Mean values of the membership results obtained by $\mathcal{A}_{\mathrm{MMSC}}$ (dotted), $\widetilde{\mathcal{A}}_{\mathrm{MMSC}}$ (dashed), and the $\mathrm{LP}_{\mathrm{MMSC}}$ solution with 2 (a), 5 (b), and 10 (c) base stations per unit disk.

unit disk, Figure 2(a) shows the mean membership values over 200 networks—for each simulated client density—for the results computed by $\mathcal{A}_{\mathrm{MMSC}}$, $\widetilde{\mathcal{A}}_{\mathrm{MMSC}}$, and the values obtained by solving $\mathrm{LP}_{\mathrm{MMSC}}$. The results depict that for this relatively low base-station density all measured values are comparable and increase with growing client density. In contrast, for a higher base-station density of 5 base stations per unit disk (cf. Figure 2(b)), a gap opens between the $\mathcal{A}_{\mathrm{MMSC}}$ and $\mathrm{LP}_{\mathrm{MMSC}}$ results. Whereas the ratio between these two result series—as mentioned before, an upper bound for the approximation ratio—rises sharply for low client densities, its increase diminishes for higher client densities, which corresponds to the $\mathrm{O}(\log n)$ approximation factor described in the theoretical analysis. Additionally, it can be observed that $\widetilde{\mathcal{A}}_{\mathrm{MMSC}}$ performs significantly better than $\mathcal{A}_{\mathrm{MMSC}}$. The reason for this effect lies in the fact that $\mathcal{A}_{\mathrm{MMSC}}$ multiplies the $\underline{x}'$ values resulting from $\mathrm{LP}_{\mathrm{MMSC}}$ with the factor $\log n$ to obtain the probabilities employed for randomized rounding, whereas this multiplication is not performed by $\widetilde{\mathcal{A}}_{\mathrm{MMSC}}$. The approximation gap becomes even wider for higher base-station densities, such as 10 base stations per unit disk (Figure 2(c)). Our simulations showed however that beyond this base-station density no significant changes in the membership results can be observed.

The increasing gap between the simulated algorithms and the $\mathrm{LP}_{\mathrm{MMSC}}$ solution with growing base-station density can be explained by the following observation: For low base-station densities—where problem instances contain a small number of sets—a relatively large number of clients are covered by only one set, which consequently will have to be chosen in both the $\mathrm{LP}_{\mathrm{MMSC}}$ and the algorithm solutions; for high base-station densities, in contrast, the solution weights $\underline{x}'$ computed by $\mathrm{LP}_{\mathrm{MMSC}}$ can be distributed more evenly among the relatively high number of available sets, and the potential of "committing an error" during randomized rounding increases.

In summary, the simulations show that the con-

9

sidered algorithms approximate the optimum solution well on practical networks. Comparing $\mathcal{A}_{\mathrm{MMSC}}$ and $\widetilde{\mathcal{A}}_{\mathrm{MMSC}}$, it can be observed that, in practice, the latter algorithm performs even better than the former.

7 Conclusion

Interference reduction in cellular networks is studied in this paper by means of formalization with the *Minimum Membership Set Cover* problem. To the best of our knowledge this combinatorial optimization problem has 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 $NP \subset TIME(n^{O(\log \log n)})$. In a second part of the paper this lower bound is shown to be asymptotically matched by a randomized algorithm making use of linear programming relaxation techniques. The third part of the paper discusses the behavior of the algorithm on practical networks. In particular, it shows that the algorithm can be modified to perform well not only in theory but also in practice.

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. *Linear Programming*. W. H. Freeman and Company, 1983.

[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] 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.

[5] D. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.

[6] N. Karmarkar. A New Polynomial-Time Algorithm for Linear Programming. *Combinatorica 4*, pages 373–395, 1984.

[7] C. Lund and M. Yannakakis. On the Hardness of Approximating Minimization Problems. *Journal of the ACM*, 41(5):960–981, 1994.

[8] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.

[9] 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.

[10] P. Raghavan and C. Thompson. Randomized Rounding: A Technique for Provably Good Algorithms and Algorithmic Proofs. *Combinatorica*, 7(4):365–374, 1987.

[11] A. Srinivasan. Improved Approximations of Packing and Covering Problems. In *Proc. of the 27th ACM Symposium on Theory of Computing*, pages 268–276, 1995.