The Complexity of Base Station Positioning in Cellular Networks*

Christian Glaßer, Steffen Reith, and Heribert Vollmer

Theoretische Informatik
Universität Würzburg
Am Hubland
D-97074 Würzburg, Germany

{glasser|streit|vollmer}@informatik.uni-wuerzburg.de

Abstract

We consider two optimization problems for cellular telephone networks, that arise in a recently discussed ITU proposal for a traffic load model. These problems address the positioning of base stations (on given possible locations) with the aim to maximize the number of supplied demand nodes and minimize the number of stations that have to be built. We show that these problems are hard to approximate, but their Euclidean versions allow a polynomial-time approximation scheme (PTAS). Furthermore, we consider other related optimization problems.

1 Introduction

The *demand node concept* has been used recently as a simplified traffic load model for cellular networks, see, e.g., [TGT96, Tut98]. A proposal for a standardization of this model is currently discussed by the International Telecommunication Union [ITU99]. In this model, so called *demand nodes* are used to describe the arising phone traffic quantitatively. A demand node represents the center of an area with a certain teletraffic demand. Each node stands for the same portion of traffic load. Densely populated areas, hence, will lead to large numbers of demand nodes, while rural regions have a sparse distribution of nodes.

One of the main problems to be solved during a network design stage is to solve the question how to position the base stations. Typically, one has to select a base station configuration from a set of possible configurations such that an optimum quality of service is achieved, i.e., as much as possible of the traffic is served while interferences have to be kept to a minimum. In the demand node model, this will usually mean that nearly all demand nodes have to be *supplied*, where the meaning of "supplied" may depend on a number of parameters.

We believe that while standardization of the demand node concept is still discussed, it is important not only to have simulation and planning tools [TLT97] but also to have precise results about (non-) solvability and (non-) approximability of combinatorial problems that arise in this model. In the present paper we consider the following optimization problems:

Maximize number of totally supplied nodes **(**MTSN**) [MN99]:**
Given is a set N of demand nodes (DN), a set B of potential locations for base stations (BS),

*Part of this work has been supported by NORTEL External Research.

a budget $k \in \mathbb{N}$ (i.e., a maximum of k base stations is allowed to be built) and an interference factor $\gamma \in \mathbb{Q}$ with $\gamma > 1$. Furthermore, for every $i \in B$ and $j \in N$, variables $u_{i,j} \in \mathbb{Q}_+$ and $t_{i,j} \in \{0, 1\}$ are given with the following meaning: The variable $t_{i,j}$ expresses if a demand node is supplied by a base station; more precisely, DN j is supplied by BS i if and only if $t_{i,j} = 1$. The variable $u_{i,j}$ gives us the intensity of the signal of BS i in DN j (if we would build this station).

If DN j is supplied by BS i, we say that it is *totally supplied* by BS i if and only if the corresponding signal in DN j is essentially stronger than the sum of all other (interfering) signals in this demand node, formally:

$$u_{i,j} > \gamma \cdot \sum_{\substack{\text{built BS } \mu \\ \mu \neq i}} u_{\mu,j}$$

Now the aim is to find a collection of at most k base stations, such that the number of totally supplied demand nodes is maximal.

Minimize number of base stations **(**MBS**):**
Here we have almost the same situation as in the problem MTSN, with the difference that a number ℓ is given instead of the budget k. The task is to find a minimum number of base stations such that at least ℓ demand nodes are totally supplied.

Hence, in both problems we consider interferences with signals from other base stations as they appear, e.g., in the present TDMA-based GSM systems as well as in the IS-95 system or in future UMTS networks that rely on the CDMA technology.

It is known that the above problems are (as decision problems) NP-complete. We first show that their optimization versions are even hard to approximate (improving a result for MTSN given in [MN99]). Our hardness result builds on the recent characterization of the class NP by so called probabilistically checkable proofs (PCP), leading to strong non-approximability theorems for the CLIQUE problem (see [AL97]). We show how CLIQUE reduces to our problems in such a way that the approximation ratio is preserved. This shows that to compute approximate solutions for MTSN and MBS is infeasible; hence in order to be able to construct efficient algorithms for practical applications, we have to restrict the model in an appropriate way. In this paper, we consider the Euclidean versions of MTSN and MBS, defined as follows:

By EMTSN (EMBS) let us denote the subproblem of MTSN (MBS, resp.) where we make the following assumptions:

- The demand nodes and the base stations are located in the Euclidean plane.
- The activity of the base stations is local (i.e. the intensity of the signal is 0 outside a range of action D).
- There exists a minimal distance d between different base stations that have to be built.

We want to stress that in reality, all these assumptions are certainly given; hence these subproblems still capture the needs of practical network optimization.

A result of this paper is that the Euclidean problems have polynomial-time approximation schemes, i.e., for any given approximation ratio r, there is a polynomial-time approximation algorithm. The algorithms we construct use a refinement of a technique developed by Hochbaum and Maas in [HM85], the so called *shifting strategy*. This technique is applicable to a lot of geometric covering and packing problems. A different approach to approximate hard geometric optimization problems was recently developed by Arora [Aro98].

In Sect. 5, we additionally define and examine two problems where by definition we neglect interferences. These problems concern the planning of base station configurations, not the selection of a particular (optimal) configuration out of a set of possible configurations. The algorithms we obtain are useful in a first stage of approximate planning of a cellular network.

2 Non-Approximability Results

We consider undirected graphs without self-loops. A clique is a subset C of nodes of a given graph such that every two different vertices in C are joined by an edge. The maximization problem of finding a clique as large as possible in a given graph is denoted by CLIQUE (of course the cost of a clique C is its cardinality). The size of a given CLIQUE-instance is the number of nodes of the corresponding graph, the size of an MTSN-instance is the number of demand nodes.

Let $h(n) \geq 1$. We say that an algorithm A is an $h(n)$-approximation algorithm for a maximization problem Π, if the solution produced by A for an input of size n is at least the optimal solution divided by $h(n)$. We say that an algorithm A is an $h(n)$-approximation algorithm for a minimization problem Π, if the solution produced by A for an input of size n is at most the optimal solution times $h(n)$. In both cases, $h(n)$ is called the *approximation ratio* [ACG+99].

Theorem 2.1 *Let $h : \mathbb{N} \to \mathbb{N}$. If there exists a polynomial-time $h(n)$-approximation algorithm A for* MTSN, *then there is also a polynomial-time $h(n)$-approximation algorithm A' for* CLIQUE.

Proof. The algorithm A' works as follows. On input of a graph G with n nodes, the algorithm generates an MTSN instance I which consists of n potential locations for base stations and n corresponding demand nodes. I has the interference factor $\gamma = 2$ and the budget $k = n$, i.e., all possible base stations are allowed to be built. A demand node i is supplied by base station j (i.e., $t_{i,j} = 1$) if and only if $i = j$. We introduce three intensities of signals: a *weak signal* of intensity $1/(8n^2)$, a *medium signal* of intensity $1/(3n)$ and a *strong signal* of intensity $1/(2n)$. Demand node j receives a weak (medium) signal from base station i if and only if $i \neq j$ and the nodes i and j are joined (are not joined, resp.). Demand node j receives a strong signal from base station i if and only if $i = j$. Now A' simulates the algorithm A on input I and obtains a collection B of (at most $k = n$) base stations. Finally A' determines the set D of demand nodes which are totally supplied by the collection B and returns this set D.

The idea behind the construction of I is as follows. With each node i in graph G we associate the pair (base station i, demand node i). A weak signal can be considered as a negligible signal (we choose an intensity near to 0, since 0 is not allowed), i.e., it can neither supply the demand node nor prevent (by interference) that the demand node is supplied totally. A medium signal is too weak to supply the demand node, but it is strong enough to prevent that the demand node is supplied totally. A strong signal can supply the demand node and it prevents that the demand node is supplied by some other signal. Note that demand node i can only be supplied by base station i. In this construction, a clique will be represented by a set of totally supplied demand nodes.

Let G be a graph with n nodes, and let C be a maximal clique in G. We have to show that A' returns a clique of size $\geq |C|/h(n)$ on input G. First of all observe that the size of G and the size of I are equal to n. We want to estimate the number of demand nodes which are totally supplied if we build all base stations $i \in C$ (and only these). Note that demand node i is supplied by base station i, for all $i \in C$. Moreover, demand node j receives a weak signal from base station i for different $i, j \in C$, since these nodes are joined. It follows that

each demand node $i \in C$ is totally supplied. Therefore, the number of demand nodes which are totally supplied in an optimal solution of I is greater than or equal to $|C|$. Thus A returns a collection B of base stations which supplies $m' \geq |C|/h(n)$ demand nodes totally. We denote this set of totally supplied demand nodes by D. Since each demand node $i \in D$ is supplied, we obtain $D \subseteq B$. Since each demand node $j \in D$ is totally supplied, this demand node receives weak signals from all base stations $i \in B \setminus \{j\}$. It follows that the nodes i and j are joined for all different $i, j \in D$. Thus with D we have found a clique of size $m' \geq |C|/h(n)$. It follows that A' is a polynomial-time $h(n)$-approximation algorithm for CLIQUE. □

Hence we showed that the problem CLIQUE reduces to MTSN. However, CLIQUE cannot be well-approximated; more precisely, it is known that there is no $(n^{1-\varepsilon})$-approximation algorithm for CLIQUE, unless every set in NP has a polynomial-time probabilistic zero-error algorithm (i.e., NP = ZPP), and that there is no $(n^{0.5-\varepsilon})$-approximation algorithm for CLIQUE, unless every set in NP has a polynomial-time algorithm (i.e., NP = P) [Hås99].

Corollary 2.1 *Let $\varepsilon > 0$.*

1. MTSN *has no polynomial-time $(n^{1-\varepsilon})$-approximation algorithm, unless every set in* NP *has a polynomial-time Las Vegas algorithm* (i.e., NP = ZPP).
2. MTSN *has no polynomial-time $(n^{0.5-\varepsilon})$-approximation algorithm, unless* NP = P.

Hence for each $0 \leq \sigma < 1$ it is very unlikely that there is a polynomial-time algorithm computing solutions for MTSN-instances such that, for each input of size n, the number of totally supplied demand nodes is greater than or equal to the optimal number divided by n^σ. Our result improves a hardness result from [MN99], where (using a reduction from the dominating set problem) only a weaker non-approximability statement was given.

To show the non-approximability of the MBS problem, we use the minimization problem *exact cover* (ECOV for short) which is defined as follows. Given a set $I = \{1, \ldots, n\}$ and a list $S_1, \ldots, S_m$ of subsets of I, we have to find a minimal collection $C \subseteq \{1, \ldots, m\}$ of list elements which cover I exactly, i.e., $I = \bigcup_{i \in C} S_i$ and $S_i \cap S_j = \emptyset$ for different $i, j \in C$. The size of the ECOV-instance is $|I|$, the cost of the solution $|C|$ and the size of an MBS-instance is the number of demand nodes.

Theorem 2.2 *Let $h : \mathbb{N} \to \mathbb{N}$. If there exists a polynomial-time $h(n)$-approximation algorithm A for* MBS, *then there is also a polynomial-time $h(n)$-approximation algorithm A' for* ECOV.

Proof. On input of an ECOV-instance $I = \{1, \ldots, n\}$ with $S_1, \ldots, S_m \subseteq I$, the algorithm A' constructs the following MBS-instance. The number of base stations is m, the number of demand nodes is n, the interference factor is $\gamma = 2$, and we require a solution where *all* demand nodes are totally supplied, i.e., $l = n$. In this MBS-instance only two different intensities of signals will appear, namely a *weak signal* of intensity $1/4m^2$ and a *strong signal* of intensity $1/2m$. More precisely, demand node j receives a strong signal from base station i if and only if $j \in S_i$. Demand node j is supplied by base station i if and only if it receives a strong signal from this station. In this construction a base station j represents the set S_j, since it supplies exactly the demand nodes $i \in S_j$ with strong signals. Now A' simulates the algorithm A on the constructed MBS-instance, obtains a set of base stations that have to be built and returns exactly this set.

To show that A' is an $h(n)$-approximation algorithm for ECOV, we assume that $\tilde{C}$ is a minimal solution of the given ECOV-instance. It follows that each demand node j is supplied

by exactly one base station $i \in \tilde{C}$, since I has an exact cover by the sets S_i with $i \in \tilde{C}$. Thus all demand nodes are totally supplied if we build the base stations from $\tilde{C}$. This shows that the minimal solution of the constructed MBS-instance is at most $|\tilde{C}|$. Therefore, A returns a solution C for the MBS-instance such that $|C| \leq |\tilde{C}| \cdot h(n)$ (note that the MBS-instance and the ECOV-instance are of size n).

It remains to show that C induces an exact cover of I. Observe that each demand node j is supplied by exactly one base station from C, otherwise the interferences would prevent this demand node to be totally supplied. By the construction of the MBS-instance this means that for every $j \in I$ there is exactly one $i \in C$ such that $j \in S_i$. Hence I has an exact cover by the sets S_i with $i \in C$. □

From [Fei96] it is known that MBS cannot be well-approximated; more precisely, it is shown that there is no $((1-\varepsilon)\ln n)$-approximation algorithm, unless every set in NP has a quasipolynomial-time algorithm (i.e., a deterministic algorithm running in time $2^{\log^d n}$ for some d). So we obtain the following consequence.

Corollary 2.2 *Let $\varepsilon > 0$.* MBS *has no polynomial-time $((1-\varepsilon)\ln n)$-approximation algorithm, unless every set in* NP *has a quasipolynomial-time algorithm (i.e., a deterministic algorithm running in time $2^{\log^d n}$ for some d).*

Hence for each $0 \leq \sigma < 1$ it is very unlikely that there is a polynomial-time algorithm computing solutions for MBS-instances such that, for each input of size n, a solution is found where the number of base stations is at most $\sigma \cdot \ln n$ times of the optimal number.

All in all we see that the problems MTSN and MBS cannot be approximated efficiently. But experience from reality shows that we can make additional assumptions about occuring problem instances. In the next section, we therefore consider the Euclidean versions of both problems, where we additionally assume a maximal radius of activity for each base station, and a minimum distance between different base stations.

3 A Polynomial-Time Approximation Scheme for EMTSN

Say that a problem Π admits a polynomial-time approximation scheme (PTAS for short), if for every rational number $r > 1$ there is a polynomial-time algorithm A_r that r-approximates Π. The following algorithm constitutes such an approximation scheme for EMTSN.

Given are the sets N of demand nodes, B of base station locations, budget k, interference factor γ and variables $t_{i,j}$, $u_{i,j}$ for every $i \in B$, $j \in N$. Additionally we assume that we have a fixed range of action D and minimal distance d between base stations.

Let the set N of the given demand nodes in the plane be enclosed in the rectangular area A. We fix an integer $l > 0$, the so called shifting parameter. This number will later determine the approximation ratio of our algorithm. By increasing l we will get better approximations.

1. Fix a division of A into horizontal and vertical strips of width D, where each strip is left (up) closed and right (down) open. (This means that for every strip, the points of the left (up) borderline belong to this strip, while the points on the right (down) borderline do not belong to it.) By considering groups of l consecutive strips, we obtain a grid of width lD, i.e., a partition of the plane A into squares of size $lD \times lD$. Such a partition can be shifted down (to the right) by $D, 2D, 3D, \ldots, (l-1)D$ before we obtain the same partition again.

Hence there are l^2 ways of partitioning A into grids of width lD. We denote the resulting partitions (we will call them *shift partitions*) by $S_1, \ldots, S_{l^2}$. Notice that such a partition can be described by the coordinates of a single point in the plane. Therefore, we can fix a partitioning of A in polynomial time even if area A is large.

2. Each partition S_i consists of several squares of size $(lD)^2$. Notice that the number of such squares is possibly superpolynomial in the input size, but in this case almost all squares do not contain any demand node. Therefore, it suffices to consider the small number (polynomially bounded in the input size) of non-empty squares, i.e., squares containing at least one demand node. This can be done efficiently because it is possible to keep all demand nodes in a suitable data structure, such that the corresponding square can be easily determined from the position of the demand node in the Euclidean plane. In each such square s we delete all demand nodes that are located in the border-strips of width D; thus we obtain a new square s' of the same size but with no demand nodes located in distance less than D to the border. For $0 \leq j \leq |B|$ we want to solve the local maximization problem in s' with budget j, i.e., we maximize the number of totally supplied demand nodes in s', where we are allowed to build up to j base stations in s'. We can do this by exhaustive search, because we can build at most $c =_{df} 4(lD)^2/(\pi d^2)$ base stations; this follows by the simple fact that circles of radius $d/2$ around every base station are not allowed to intersect because of the minimal distance d between two base stations. A more careful analysis shows that even $2(lD)^2/(d^2\sqrt{3})$ is an upper bound on the number of base stations per square; however, the only important point is that c is a constant. Thus, for each square s, we obtain a $(c+1)$-dimensional vector v_s, where the j-th component is the maximal number of totally supplied demand nodes in s' if we are allowed to build at most j base stations.

3. Now we have reached another optimization problem, which is a kind of knapsack-problem. We are given a list L of $(c+1)$-dimensional vectors $v_s = (v_{s,0}, \ldots, v_{s,c})$ of natural numbers (one vector for each non-empty square s) and a maximal weight k (our budget). We have to find a total mapping $f\colon \{1, \ldots, |L|\} \to \{0, \ldots, c\}$ giving us for each square, which contains at least one demand node, the number of base stations to be built such that $\sum_{1 \leq i \leq |L|} f(i) \leq k$ and $\sum_{1 \leq i \leq |L|} v_{i,f(i)}$ is maximal. In other words, we want to distribute at most k base stations to $|L|$ squares, such that the number of totally supplied demand nodes is maximal. Since the $v_{i,j}$ are small numbers, we can solve this problem in polynomial time with the technique of dynamic programming as we describe next (this technique is a slight modification of an algorithm for the knapsack problem which can be found in [Pap94, pp. 203ff]).

   Let
   $$W(j,v) = \min\Big\{ \sum_{1 \leq i \leq j} f(i) \;\Big|\; f\colon \{1,\ldots,j\} \to \{0,\ldots,c\} \text{ and } \sum_{1 \leq i \leq j} v_{i,f(i)} = v \Big\}$$
   for $0 \leq v \leq |N|$ and $1 \leq j \leq |L|$.

   So the function $W(j,v)$ gives us the minimal number of base stations located in the first j squares such that exactly v demand nodes can be supplied totally.

   It is easy to see that the following recursive procedure determines these values: Let $W(0,v) = \infty$ for all $0 \leq v \leq |N|$. For $0 \leq v \leq |N|$ and $0 \leq j < |L|$, we set
   $$W(j+1,v) = \min\{\, W(j, v - v_{j+1,k}) + k \mid 0 \leq k \leq c \,\}.$$

   Therefore, we can compute all values $W(k,v)$ in time $O(|N| \cdot |L|)$.

Now we can search the largest $v_{\max}$ such that $W(|L|, v_{\max}) \leq k$; this is the solution of the knapsack-problem.

4. For the base station problem this means that we can supply $v_{\max}$ demand nodes totally if we build at most k base stations. Hence each partition S_i induces a number of demand nodes being totally supplied by at most k base stations. Under all these solutions we choose one having the largest number of totally supplied demand nodes.

Now we want to show that the relative error of this algorithm is at most $4/l$. We fix an optimal solution of the given problem instance I. On the one hand, for each demand node there exist exactly $4l-4$ partitions in which this demand node is located on a border-strip. On the other hand there are l^2 partitions at all. By the pigeon hole principle we obtain that there exists a partition S' such that the following holds: At most a fraction of $(4l-4)/l^2 \leq 4/l$ of the totally supplied demand nodes (of the optimal solution) are located on a border-strip. If we delete all demand nodes on all border-strips of S' we obtain a new instance I' of our problem. It is easy to see that the solution for S' obtained by local optimization (first optimize each single square, then solve the knapsack-problem) equals the solution for I' obtained by global optimization, because we have deleted all demand nodes located outside a square and inside the range of action of a base station in this square. Thus base stations in one square cannot influence demand nodes in another square. Therefore, if we carry out this local optimization, we find in polynomial time an optimal solution for problem instance I'.

From the optimal solution for I we obtain a (not necessarily optimal) solution for I' if we hide the (small amount of) totally supplied demand nodes located on the border-strips of S'. This gives $\mathrm{OPT}_I \leq \mathrm{OPT}_{I'}/(1-4/l)$, which shows that our approximated solution is near to the optimal solution. Since every optimal solution for I' induces a solution for I, we have:

$$\mathrm{OPT}_{I'} \leq \mathrm{OPT}_I \leq \mathrm{OPT}_{I'}/(1-4/l)$$

Therefore, at least for partition S' our algorithm finds a solution having a relative error $\leq 4/l$, hence the overall solution of the algorithm will not have a larger error.

Since the l in the above argument is a number which we can choose beforehand, we see that for every approximation ratio $r > 1$ we obtain a polynomial algorithm for the EMTSN problem. Hence we proved the following theorem:

Theorem 3.1 *The problem* EMTSN *admits a PTAS.*

4 Approximating EMBS in Polynomial Time

Let us now turn to the EMBS problem. We take over the notations from the problem EMTSN, with the difference that ℓ, the minimal number of totally supplied demand nodes, is given (instead of budget k). Since EMBS is a minimization problem, where we demand a certain amount of success (at least ℓ totally supplied demand nodes), it is possible that there exist problem instances which do not have solutions. Even if an instance has a solution, it could happen that this is the only one, and in this case, every approximation algorithm has to find this optimal solution. So we cannot hope to find an approximation algorithm for this strong problem, and we have to weaken the requirement to supply at least ℓ demand nodes. Instead, we proceed as follows:

Starting from the PTAS for EMTSN, for every ε we construct an algorithm A_ε such that the following holds: For every instance I, which asks for ℓ totally supplied demand nodes and which has an optimal solution of b base station, our algorithm gives a solution with $b' \leq b$ base stations and $\ell' \geq \ell/(1+\varepsilon)$ totally supplied demand nodes.

We now describe the construction of algorithm A_ε: Let I be an instance with a set of base stations B and a set of demand nodes N, such that b is the minimal number of base stations which have to be built if we want to supply at least ℓ demand nodes totally. Let B_ε be the approximation algorithm for EMTSN whose solutions have a relative error $\leq \varepsilon$. On input I, algorithm A_ε applies B_ε on instance I with budget $k = 1, \ldots, |B|$ and obtains a list of numbers of totally supplied demand nodes $\ell_1, \ldots, \ell_{|B|}$. Let i be the smallest number with $1 \leq i \leq |B|$ such that $\ell_i \geq \ell/(1+\varepsilon)$. It follows that $i \leq b$. Therefore, on instance I with budget i, algorithm B_ε outputs a collection of $\leq b$ base station which supply $\ell' \geq \ell/(1+\varepsilon)$ demand nodes totally. Our algorithm A_ε outputs exactly this solution.

The polynomial-time approximation algorithm above leads to the following theorem:

Theorem 4.1 *Let $\varepsilon > 0$. There exists an algorithm that, for any instance I of* EMBS *where b is the minimal number of base stations required to supply ℓ demand nodes, computes a solution with at most b base stations and $\ell' \geq \ell/(1+\varepsilon)$ totally supplied demand nodes.*

5 Optimizing Networks without Interferences

5.1 A Polynomial-Time Approximation Scheme for Base Station Localization

We define the following geometric optimization problem:

Minimize number of base station locations **(EBSL$_{NI}$):**
Given is a set N of demand nodes located in the Euclidean plane. The task is to compute a minimum cardinality set of locations for base stations such that all demand nodes are supplied; where here supplied means that the demand node is within distance D from a built base station.

In this problem we hence do not fix a set of potential locations for the base stations; the goal consists in computing these potential locations in an optimal way. On the other hand, let us clearly point out that we do not take interferences of base station signals into account.

We next describe an approximation algorithm for the EBSL$_{NI}$ problem. The algorithm given below is an immediate translation of the algorithm for the disk covering problem presented in [HM85] to case of the Euclidean plane. Let l be any number.

1. Let the demand nodes be enclosed in the plane in area A. As in the algorithm for EMTSN, construct shift partitions $S_1, \ldots, S_{l^2}$, where each partition divides the plane into a grid of squares of size $lD \times lD$.
2. For each partition S_i, locally solve the maximization problem for each non-empty square s in S_i. The solution for S_i then is obtained as the union of all base stations over all non-empty squares s.
3. Under all solutions for the different S_i, choose the minimum.

This algorithm works in polynomial time, and solves the EBSL$_{NI}$ problem with approximation ratio at most $(1+1/l)^2$, see [HM85]. Thus we conclude:

Theorem 5.1 *The problem* EBSL$_{NI}$ *admits a PTAS.*

5.2 A PTAS for EMBS_{NI}

Finally, let us again consider the problem EMBS, but this time neglecting interferences of signals from different base stations. In Sect. 4 we were not able to present a PTAS, due to the particular structure of EMBS, where we want to perform a minimization but it can well be that a solution does not exist. When we do not consider interferences, we are guaranteed that always some solution exists, and for this problem we are able to construct a PTAS.

Minimize number of base stations without interferences **(EMBS_{NI}):**
Given is a set N of demand nodes, a set B of potential locations of base stations, and a number ℓ. The task is to find a minimum number of base stations such that at least ℓ demand nodes are supplied, where here, supplied simply means that the demand nodes are within distance D from a built base station.

Note that here, in contrast to the previous subsection, potential base station locations are given. Again, we present a polynomial-time approximation scheme.

1. Let the demand nodes be enclosed in the plane in area A. As in the algorithm for EMTSN, construct shift partitions $S_1, \ldots S_{l^2}$, where each partition divides the plane into a grid of squares of size $lD \times lD$.

2. Fix a partition S_i. We locally solve the following maximization problems for each square s in S_i by exhaustive search: For each number j compute the maximal number $v_{s,j}$ of demand nodes in s that can be supplied, when we are allowed to build j base stations. Here, base stations can be build either within s or within the border strips of neighboring squares. Let us denote by $U(s)$ the square s plus these border strips, hence $U(s)$ is a square of size $(l+2) \times (l+2)$. Hence, the task is to supply a maximal number of nodes within s by building j base stations in $U(s)$.

3. Now we have a list L of $(c+1)$-dimensional vectors $v_s = (v_{s,0}, \ldots, v_{s,c})$ of natural numbers (one vector for each square s; c is a constant bounding the number of base stations that can be built in a square $U(s)$, see the description above of the algorithm for EMTSN). The entry $v_{s,j}$ gives us the number of nodes in s that can be supplied by building j base stations in $U(s)$. The problem we have to solve now is to choose for each square s a number $v_{s,j}$ such that the sum of all these is at least ℓ, while the number of base stations that have to be build (i.e., the sum of the j's over all picked $v_{s,j}$) is minimal. This is a knapsack problem which we solve using a dynamic programming approach, analogous to the algorithm for EMTSN.

4. Finally, among the solutions for the partitions $S_1, \ldots, S_{l^2}$ we output the one with the minimal solution.

Let us analyze the error of this algorithm. Consider a problem instance I and an optimal solution S for I (S is a set of base stations). Consider a square s of some partition, where in the optimal solution a number of v_s demand nodes is supplied. Suppose that for this, m_s base stations have to be built in s and m'_s base stations have to be built in $U(s) - s$. Local optimization in s will therefore not lead to more than $m_s + m'_s$ base stations built in $U(s)$ in order to supply v_s demand nodes. There are l^2 partitions; over all partitions, each base station can be located at most $4l$ times in a border strip, hence, by the pigeon hole principle, there must be a partition where, the ratio of stations from S in border strips is at most $4/l$. By the above, to supply v_s nodes in square s the local optimization leads to at most $m_s + m'_s$ base stations. The sum of all m_s over all squares s is OPT_I, and there must be some partition for which the

sum of all m'_s is at most $4/l \cdot \mathrm{OPT}_I$. The solution found in this partition can therefore be not larger than $\mathrm{OPT}_I(1 + 4/l)$.

Since l can be an arbitrary number, this algorithm gives us the following theorem:

Theorem 5.2 *The problem* EMBS_{NI} *admits a PTAS.*

6 Conclusion

We presented polynomial-time approximation schemes (PTAS) for the optimization problems EMTSN, EBSL_{NI} and EMBS_{NI}. Furthermore we gave a similar family of algorithms for EMBS. This means that for any approximation ratio $r > 1$ we wish, there is a polynomial-time approximation algorithm for all four problems. One might ask if there is a uniform approximation scheme, working for all ratios r with a runtime which is polynomial in the input length and $1/(r-1)$. This would mean that the problems have a so called fully polynomial-time approximation scheme (FPTAS), see [ACG+99]. However, for all our Euclidean problems such stronger schemes cannot exist (unless P = NP). This is due to the fact that the corresponding decision problems are strongly NP-complete, and optimization problems for such problems cannot have a FPTAS [GJ79, Theorem 6.8]. Hence, approximation schemes as presented here are the best we can hope for.

It should be remarked that the algorithm we gave above for the EMTSN problem can in fact be generalized: Our algorithm maximizes the number of *totally supplied nodes*. In fact, "totally supplied" here can be replaced by any other, even much stronger, condition — let us call it *super-supplied*. As long as it can be checked in polynomial time whether a node is super-supplied (by a given collection of base stations), our algorithm will produce an approximate solution in polynomial time. This allows us, e.g., to take peculiarities of the landscape (highways, skyscrapers, etc.) into account. Note that we used the fact that the rectangular area A is embedded in the Euclidean plane *only* to determine the positions of base stations and demand nodes. None of the algorithms above makes any assumption about a metric for the field strength/interferences.

A further generalization concerns the question of *incremental* versions of the above examined problems. In this case, a problem instance consists, besides those components specified above, also of a description of an existing network of some already built base stations. The task now is to extend this network by building up to k new base stations (for EMTSN), or by building a minimal number of additional stations in order to supply a number of at least ℓ demand nodes (for EMBS). It can be observed that our algorithms still work in this context, hence we can state the same complexity results as above in Theorems 3.1 and 4.1.

Acknowledgment. We are grateful to Matthias "Lotti" Galota, Kenji Leibnitz, Phuoc Tran-Gia, Kurt Tutschku, and Klaus W. Wagner for helpful discussions. In particular we thank Rudolf Mathar for initiating this work by providing us with a formal definition of the MTSN problem.

References

[ACG+99] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation – Combinatorial Optimization Problems and Their Approximability.* Springer Verlag, Berlin Heidelberg, 1999.

[AL97] S. Arora and C. Lund. Hardness of approximations. In D. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*, chapter 10, pages 399–446. PWS Publishing Company, Boston, 1997.

[Aro98] S. Arora. Polynomial-time approximation schemes for Euclidian TSP and other geometric problems. *Journal of the ACM*, 45:753–782, 1998.

[Fei96] U. Feige. A threshold of $\ln n$ for approximating set cover. In *Proceedings 28th Symposium on Theory of Computing*, pages 314–318. ACM Press, 1996.

[GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness.* Freeman, New York, 1979.

[Hås99] J. Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182:105–142, 1999.

[HM85] D. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM*, 32:130–136, 1985.

[Hoc97] D. Hochbaum, editor. *Approximation Algorithms for NP-hard Problems.* PWS Publishing Company, Boston, 1997.

[ITU99] ITU-T Focus Group on Traffic Engineering for Personal Communication. E.750-Series of Recommendations on Traffic Engineering Aspects of Networks Supporting Mobile and UPT Services. Technical reference, The International Telecommunication Union, Geneva, CH, 1999.

[MN99] R. Mathar and T. Niessen. Optimum positioning of base stations for cellular radio networks. In *Wireless Networks*, 1999.

[Pap94] C. H. Papadimitriou. *Computational Complexity.* Addison-Wesley, Reading, MA, 1994.

[TGT96] K. Tutschku, N. Gerlich, and P. Tran-Gia. An integrated approach to cellular network planning. In *Proceedings 7th International Network Planning Symposium*, Sydney, 1996.

[TLT97] K. Tutschku, K. Leibnitz, and P. Tran-Gia. ICEPT – an integrated cellular network planning tool. In *Proceedings IEEE/VTS 47th Vehicular Technology Conference*, Phoenix, AZ, 1997.

[Tut98] K. Tutschku. Demand-based radio network planning of cellular communication systems. In *Proceedings IEEE Infocom*, San Francisco, CA, 1998.