A Polynomial-Time Approximation Scheme for Base Station Positioning in UMTS Networks*

Matthias Galota, Christian Glaßer, Steffen Reith, Heribert Vollmer
Theoretische Informatik, Universität Würzburg
Am Hubland
D-97074 Würzburg, Germany
[galota, glasser, streit, vollmer]@informatik.uni-wuerzburg.de

ABSTRACT

We consider the following optimization problem for UMTS networks: For a specified teletraffic demand and possible base station locations, choose positions for base stations such that

- the construction costs are below a given limit,
- as much teletraffic as possible is supplied,
- the ongoing costs are minimal, and
- the intra-cell interference in the range of each base station is low.

We prove that for a particular specification of teletraffic (the so called demand node concept), this problem has a polynomial-time approximation scheme, but cannot have a fully polynomial-time approximation scheme unless P = NP.

Categories and Subject Descriptors

C.2.1 [**Computer–Communication Networks**]: Network Architecture and Design—*Wireless communication*; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Algorithms, Design, Theory

Keywords

Approximation Algorithms, Network Design

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

1. INTRODUCTION

1.1 CDMA Networks

During the planning stage for a new cellular network, one of the main tasks for a provider is to decide where to place the base stations (i.e., the antennas). In order to accomplish this, data about the expected teletraffic demand is needed. A widely accepted way of describing phone traffic quantitatively is by providing so called *demand node maps* [15]. A demand node represents the center of an area with a certain teletraffic demand. Each node stands for the same amount of traffic load. Thus, densely populated areas will lead to a large number of demand nodes while rural areas will lead to a sparse distribution. Demand node maps are produced in a canonical way from given stochastic data (a standardization for this is currently discussed by the International Telecommunication Union [9]).

When a base station is built, its signals reach all users within a certain range: We will say that the corresponding demand nodes are *supplied*. For a base station i we will denote the set of all supplied demand nodes by N_i. For network architectures based on the so called *time division multiple access* technology (TDMA), as for example the current European standard system GSM, building one base station could have an effect on other base stations. For example, neighbored base stations had to use frequencies that are sufficiently far apart (frequency assignment problem). Also, interferences between different base stations had to be considered; it could well be that a demand node is located near to a base station i, nevertheless it is not supplied (i.e., not in N_i) because it is disturbed by too strong signals from other stations.

Other network architectures, as already existing in the United States (the IS-95 network) and as currently developed in Europe (UMTS networks), rely on *code division multiple access* technology (CDMA), see [14, Sect. 8.4.1 and Sect. 10.4]. This means that signals, both from a base station and from a mobile station, are multiplied by a very large bandwidth signal called the *spreading signal*, a certain kind of pseudonoise code sequence. All mobile users and all base stations have their own pseudorandom codewords which are approximately orthogonal to all other codewords. This implies that users in such a CDMA system use the same carrier frequency and may transmit simultaneously. This leads to very high transfer rates. (In contrast to this, time slots were reserved for each user in older *time division multiple access* networks.)

For the network planning stage, this means the following: As a first step we consider the model in this paper, where different base stations are regarded as independent from each other; that is, we neglect frequency conflicts. (See the conclusions for a discussion of this simplification over reality.) Furthermore, the network has a soft capacity limit. Increasing the number of mobile users supplied by a base station (for instance by increasing the broadcast range of the base station) raises the noise floor within the cell of the station. While this will lead eventually even to technical limits (we might run out of different codewords), it already much earlier causes a loss with respect to the quality of the transmitted signals.

The above considerations lead to some objectives that have to be aimed at by a provider when planning a cellular network. We assume the provider has a demand node map of her area, and furthermore, we assume that she identified possible locations for base stations (e.g., church towers if the congregation agrees, roofs of industrial buildings if the company gets special phone discounts, roofs of public buildings, etc.). The provider now has to choose a subset of the possible locations and built base stations there. The choice will be determined by the following:

1. To reach a high coverage of the area, as many of the given demand nodes as possible have to be supplied. For each covered node, a certain gain will be obtained.
2. Usually, there is a certain maximal budget that the provider is allowed to spend during the construction.
3. Built base stations have to be supported; hence they will cause ongoing costs.
4. Since in our model we do not consider interference, it will cause no problem if a node is covered by more than one station. However, within the cell of each base station, the intra-cell interference will rise if the number of supplied users rises. Higher interference will lead to lower customer satisfaction and, hence, cause a certain loss.

Of course, the provider has to find a way to compare the different types of costs—the singular construction costs, the ongoing supporting costs and the not even monetary interference costs—to be able to charge them up against each other.

Assuming that such normalized fictional costs have been found, the profit reached by the provider is computed by adding up all these gains and losses. The goal is to maximize this profit.

1.2 Formal Definition of the Problem

Next, we present a formal definition of the problem described in Subsect. 1.1. Unless stated otherwise we stipulate that all variables are considered as natural numbers.

Let B denote a set of possible base station locations, and let N denote a set of demand nodes. In practical circumstances, the following requirements are certainly given:

(R1) The demand nodes and the possible base station locations are located in the Euclidean space. To facilitate the proofs we assume that they are even located in the two-dimensional Euclidean space, i.e., in the Euclidean plane.

(R2) Every base station $i \in B$ has a maximal broadcast range $r_{\max}(i)$. For every actually built base station i we have to determine an actual broadcast range $r(i) \leq r_{\max}(i)$. Again to simplify presentation, we assume that the signal strengths obey the Euclidean metric; this means that the sets N_i of those demand nodes that are supplied by i are then given by $N_i = \{ j \mid \mathrm{dist}(i,j) \leq r(i) \}$. We discuss non-Euclidean field strengths, for example areas with no reception, in the conclusions.

The problem BSPC is now given as follows.

***Base Station Positioning in CDMA Networks* (BSPC):**
Given are a set N of demand nodes and a set B of possible base station locations (both located in the Euclidean plane), a budget b, a gain e for supplying a demand node, construction costs $c(i)$, ongoing costs $o(i)$ for supporting a station at location $i \in B$, a cost $h(i,k)$ for the intra-cell interference that occurs when in the range of base station i a number of k demand nodes is supplied, and a maximal broadcast range $r_{\max}(i)$ for every base station $i \in B$. The aim is to find a subset $S \subseteq B$ of base stations to be built and a broadcast range $r(i) \leq r_{\max}(i)$ for every $i \in S$ such that the overall construction costs $\sum_{i \in S} c(i)$ are at most b and the following profit P is maximal:

$$
\begin{aligned}
P \quad =_{\mathrm{def}} \quad & e \cdot \overbrace{\left|\{ j \mid \exists i \in S \colon \mathrm{dist}(i,j) \leq r(i) \}\right|}^{\text{number of supplied demand nodes}} \\
& - \overbrace{\sum_{i \in S} o(i)}^{\text{ongoing costs}} \\
& - \overbrace{\sum_{i \in S} h\left(i, \left|\{ j \mid \mathrm{dist}(i,j) \leq r(i) \}\right|\right)}^{\text{cost of intra-cell interference}}
\end{aligned}
$$

In practice we can make additional assumptions:

(R3) The maximal broadcast range is bounded by a constant (currently, approximately 60 km). Also, it is clear that every base station has a certain spatial extension. This means that there exist constants d, D such that $\max\{ r_{\max}(i) \mid i \in B \}$ is bounded from above by D and the the distance between two built base stations is at least d.

In our algorithms we will consider either (R3) or the following relaxed condition:

(R3') The maximal number of actually built base stations in any circular area with radius $\max\{\, r_{\max}(i) \mid i \in B \,\}$ is bounded above by some constant C.

It is clear that, if we assume (R1), (R2), and (R3), this implies (R3') but not necessarily vice versa. Hence it is enough to solve the case (R1), (R2), and (R3').

For constants d, D we define $\mathrm{BSPC}^{d,D}$ as the restriction of BSPC to instances satisfying (R3). Similarly, for a constant C we define BSPC^{C} as the restriction of BSPC to instances satisfying (R3'). For the following complexity considerations, we define the size of an instance as the length of an encoding of all the above components, where numbers are encoded in binary.

1.3 Relation to Previous Work

Design problems for cellular networks have been examined with respect to frequency assignment problems (see, e.g., [11]) and with respect to power range assignment (see [3]). However, the problem of positioning base stations in an optimal way has not been considered in depth from a complexity-theoretic point of view. While simulation and planning tools (see [16]) and solutions based on linear programming (see, e.g., [12]) are known, not much has been proven about (non-)solvability and (non-)approximability of arising combinatorial problems. In [4], optimum non-approximability results for planning problems for telecommunication networks, where antennas are placed in balloons at a certain height, were given. In [2], an approximation algorithm for the problem of covering a plane area with cells while avoiding buildings within dangerous range of a station was given. In [6] the design problem for cellular networks that use TDMA technology was examined, e.g., GSM networks.

However, turning to third generation cellular networks that use CDMA technology, the network planning process is different from all the above. Compared to GSM networks, for each base station the interference within its cell has to be taken into account here. This posed no problem for older networks. Mainly because of this point, the problem we consider here has different combinatorial properties than the above mentioned problems. We present an approximation algorithm based on the so called *shifting strategy*, which was introduced by Hochbaum and Maass in [7]. In this technique, which is applicable to a lot of geometric problems, the plane is divided in different ways, leading to different partitions. For each partition, an approximate global optimum can be obtained by combining (exact) local optima. Then it is proved that for at least one partition the obtained solution is near to the optimal solution. Our application of the shifting strategy here is interesting, because even the solutions for particular partitions can only be obtained in an approximate way, and we need an approximation algorithm for a knapsack-like problem as a subroutine. Mainly because of this, the algorithm obtained here is more involved than the one for GSM networks given in [6].

It should be mentioned that our problem resembles a certain *facility location problem*, namely the *capacitated facility location problem with unsplittable demands*, for which an approximation algorithm was given in [10]. But there are major differences, mainly that in our problem not all demand nodes have to be supplied and the capacities of the base stations can differ.

1.4 Paper Organization

In the upcoming Sect. 2 we recall some basic notions. Sect. 3 contains the main technical part of the paper. There we present an algorithm for BSPC^{C}, witnessing that this problem has a polynomial-time approximation scheme. In Sect. 4 we prove that BSPC^{C} cannot have a fully polynomial-time approximation scheme. There, we also examine the (restricted) problem $\mathrm{BSPC}^{d,D}$ and show that even this problem does not admit an FPTAS. Finally, Sect. 5 concludes.

2. PRELIMINARIES

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* [1, p. 90].

Informally, we will say that a problem is well-approximable if it can be approximated up to any constant; formally: A problem Π admits a polynomial-time approximation scheme (PTAS for short) if, for every rational number $\rho > 1$, there is a polynomial-time algorithm A_ρ that ρ-approximates Π, see [1, p. 102]. We will prove that for every C, BSPC^{C} admits a PTAS (Theorem 1).

For problems that have a PTAS, the question of course remains if there is a uniform approximation scheme, working for all ratios ρ with a runtime which is polynomial in the input length and $1/(\rho - 1)$. This would mean that the problem has a so called fully polynomial-time approximation scheme (FPTAS), see [1, p. 111].

An optimization problem Π is NP-hard, if every language $L \in \mathrm{NP}$ reduces to Π via polynomial-time Turing-reductions, see [1, p. 30]. Let p be a polynomial; then Π_p denotes the problem obtained by restricting Π to only those instances I for which all numbers, that occur as components of input I, are bounded by $p(|I|)$. An optimization problem Π is strongly NP-hard if there is a polynomial p such that the problem Π_p is already NP-hard [1, p. 116].

3. AN APPROXIMATION SCHEME FOR BSPC^{C}

For any fixed $C \geq 0$ we want to construct an approximation scheme for BSPC^{C}. Let the given set N of 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.

Our algorithm consists of three steps.

Step 1. Fix a division of A into horizontal and vertical strips of width D, where each strip is left (up) closed and

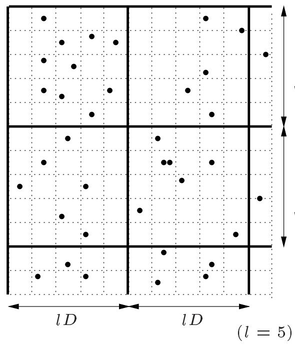

Figure 1: A partition of the plane. Dots represent base stations.

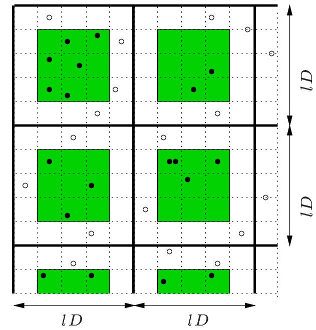

Figure 2: Clearing border strips by deleting all base stations outside the shaded squares.

right (down) open. By considering groups of l consecutive horizontal and vertical strips of width D, we obtain a partition of A into several squares of size $(lD)^2$, which form a grid of width lD; see Fig. 1. 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 (bounded by $|N|$) of non-empty squares, i.e., squares containing at least one demand node. A list $s_1, \ldots, s_n$ of these non-empty squares can be efficiently computed, since for each demand node we can determine the square containing this demand node. In each such square s_i we delete all base station locations that are in the border-strips of width D; thus we obtain a new square t_i of the same size but with no base station locations in distance less than D to the border; see Fig. 2.

Step 2. For each square t_i and each number $m \leq b$ (we call this m a *local budget*) we define a set $S(i,m)$ of base stations located in t_i such that (i) the construction cost of this set is $\leq m$, i.e., $\sum_{j \in S(i,m)} c(j) \leq m$ and (ii) if we build the base stations in $S(i,m)$ we obtain the maximal profit in t_i possible with a local budget $\leq m$ (we refer to this maximal profit as $P(i,m)$). Note that we cannot compute all the sets $S(i,m)$ and profits $P(i,m)$ for $1 \leq i \leq n$ and $m \leq b$, since i varies over an exponential range. However, a single $P(i,m)$ can be computed by exhaustive search, since in each t_i we can build at most $c = l^2 C$ base stations; this follows from our assumption that we can not build more than C base stations in any circular area with radius D, and t_i can be covered by l^2 such disks of radius D. Since the profit of building no base station is 0, the values of $P(i,m)$ are all at least 0. Moreover, it is easy to see that, for each fixed i, $P(i,m)$ as a function in m is monotonic increasing. Now we have to combine these local solutions in a nearly optimal way, which turns out as a kind of knapsack-problem. We have to distribute the overall budget b to the single squares, i.e., for each t_i we have to find a local budget m such that the sum of all these local budgets is at most b and the sum of all local profits, $P(i,m)$, is (almost) maximal.

We show that this knapsack-like problem admits a full polynomial time approximation scheme (FPTAS for short). For this it suffices to show that on input of some r and a problem instance I, we can compute in polynomial time in $|I|+r$ an approximated solution with relative error $\leq 1/r$. The proof is based on an idea of [8] (see also [13, pp. 305ff]), showing that the knapsack problem admits an FPTAS.

In the following we describe the way this approximation algorithm works. Let $p_{\max} =_{\text{def}} \max\{P(i,b) \mid i \leq n\}$, and let $P_{\max} =_{\text{def}} p_{\max} \cdot n$. Note that $P_{\max}$ is an upper bound for the obtainable profit. In the following we only consider profits that are multiples of 2^k where $k =_{\text{def}} \lfloor \log(p_{\max}/rn) \rfloor$. We will see that this is sufficient to obtain a good approximate solution. Now determine all $b_{i,p} =_{\text{def}} \inf\{m \mid m \leq b \wedge P(i,m) \geq p \cdot 2^k\}$ for $i \leq n$ and $p \leq p_{\max}/2^k$ (using binary search this is possible in time polynomial in $(n \cdot p_{\max}/2^k) \cdot \log b$. This is polynomial in $r+n+\log b$, since $p_{\max}/2^k \leq 2rn$ and $P(i,\cdot)$ is a monotone increasing function). The values $b_{i,p}$ give us complete information about the function $P(\cdot,\cdot)$ disregarding the last k bits of the profit.

Let $W(j,p) =_{\text{def}} \min\{\sum_{i=1}^{j} b_{i,p_i} \mid p_1, \ldots, p_j \in \mathbb{N}, p_1 + \cdots + p_j \geq p\}$ for $p \leq P_{\max}/2^k$ and $j \leq n$. Hence the function $W(j,p)$ gives us the minimal construction costs we have to invest into the first j squares such that the obtained profit is at least $p \cdot 2^k$ (disregarding in each square the last k bits of the profit). Since $P_{\max}/2^k \leq 2rn^2$ we can determine all values $W(j,p)$ in time polynomial in $n+r$ using a dynamic programming technique. More precisely let $W(1,p) = b_{1,p}$ for $p \leq P_{\max}/2^k$, and $W(j,p) = \inf\{\, W(j-1, p-p') + b_{j,p'} \mid p' \leq P_{\max}/2^k \,\}$ for $j \leq n$ and $p \leq P_{\max}/2^k$. Choose the largest $\bar{p}$ such that $W(n,\bar{p}) \leq b$. This means that we can reach a profit $\geq \bar{p} \cdot 2^k$ if we distribute an overall investment of $W(n,\bar{p})$ in a suitable way to the squares.

By $\bar{b}_1, \ldots, \bar{b}_n$ we denote the local budgets that are spent in squares $1, \ldots, n$ in order to obtain this profit $\bar{p} \cdot 2^k$. Next we estimate the relative error of this solution. For this we fix an optimal solution $\hat{b}_1, \ldots, \hat{b}_n$ (note that the sum of all $\hat{b}_i$ is $\leq b$).

Claim 1:

$$\bar{p} \geq \sum_{i=1}^{n} \left\lfloor \frac{P(i,\hat{b}_i)}{2^k} \right\rfloor .$$

Assume the contrary, and let $\hat{p}_i =_{\text{def}} \lfloor P(i,\hat{b}_i)/2^k \rfloor$. Note that $b_{i,\hat{p}_i} \leq \hat{b}_i$. So the sum of the local budgets $b_{i,\hat{p}_i}$ is $\leq b$

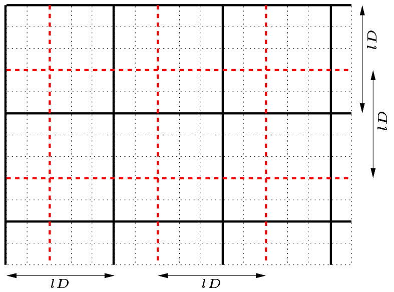

Figure 3: A second partition (dashed lines), obtained by moving the first one right by $2D$ and down by $3D$.

and the sum of the $\hat{p}_i$ is by assumption greater than $\bar{p}$. This is a contradiction to the choice of $\bar{p}$, since $\bar{p}$ was chosen to be the maximal approximate solution. This proves the claim.

Claim 2:

$$\begin{aligned}
\sum_{i=1}^{n} P(i,\hat{b}_i) &\geq \sum_{i=1}^{n} P(i,\bar{b}_i) \\
&\geq 2^k \bar{p} \\
&\geq \sum_{i=1}^{n} 2^k \left\lfloor \frac{P(i,\hat{b}_i)}{2^k} \right\rfloor \\
&\geq \sum_{i=1}^{n} P(i,\hat{b}_i) - 2^k n.
\end{aligned}$$

The first inequality holds since the $\hat{b}_i$ represent an optimal solution, the second inequality follows from the definition of $\bar{b}_1, \ldots, \bar{b}_n$, the third because of Claim 1, and the fourth follows by simple arithmetic. This proves Claim 2.

Hence, the approximate solution $2^k\bar{p}$ is at least the optimal solution minus $2^k n$. Since the value $p_{\max}$ is a lower bound for the overall optimal solution, we conclude that the approximate solution has a relative error of at most $n \cdot 2^k / p_{\max}$. Since $p_{\max}/2^k \geq rn$, this error is bounded from the above by $1/r$.

Step 3. Clearly there are l^2 different ways of partitioning A by a grid of width D into one of width lD, because each fixed partition can be shifted down (to the right) by D, $2D$, $3D$, ..., $(l-1)D$ before we obtain the same partition again; see Fig. 3.

Denote the resulting shift partitions by $P_1, \ldots, P_{l^2}$. Notice that such a partitioning 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 the area A is large. If we choose $r =_{\text{def}} l$, we obtain for every partition P_i an approximate solution p_i for the knapsack problem which has a relative error $\leq 1/l$. Under all these solutions we choose one with maximal profit.

This concludes the description of the algorithm. Next we want to show that the relative error of the whole algorithm is at most $5/l$. We fix an optimal solution of the given problem instance I. Let the profit of this solution be $\hat{p}$ and the set of built base stations be S. Now for every base station $i \in S$ let N_i be the set of demand nodes supplied by i. We assume w.l.o.g. that the sets N_i are pairwise disjoint. (If some demand node occurs in more that one of these sets, we simply delete it arbitrarily from some of the N_i in order to achieve disjointness.) Now we assign to each $i \in S$ the value $x_i =_{\text{def}} e \cdot |N_i| - o(i) - h\left(i, \left|\left\{\, k \in N \;\middle|\; \text{dist}(i,k) \leq r(i) \,\right\}\right|\right)$. If a set R of base stations is not built, then the profit lost is at most $\sum_{i \in R} x_i$. This holds because

- the number of demand nodes that are no longer supplied when we do not build base stations in R is at most $\sum_{i \in R} |N_i|$,
- the ongoing costs for supporting i is exactly $o(i)$, and
- the cost of interference we prevent by not building i is exactly

$$h\left(i, \left|\left\{\, k \in N \;\middle|\; \text{dist}(i,k) \leq r(i) \,\right\}\right|\right).$$

Note that each x_i is non-negative, since otherwise the base stations with negative x_i could be omitted and thus a solution with higher profit could be achieved.

For each partition P_j, let R_j be the set of base stations built in the optimal solution that are located in a border strip of that partition. We consider the sums $C_j =_{\text{def}} \sum_{i \in R_j} x_i$ for $1 \leq j \leq l^2$, where C_j is an upper bound for the profit we lose in partition j by deleting all base stations in the border strips.

Since for each base station $i \in S$ there exist exactly $4l-4$ partitions where i is on a border strip, each x_i occurs in exactly $4l-4$ of these C_j. Hence the sum of all C_j is equal to $\sum_{j=1}^{l^2} C_j = (4l-4) \sum_{i \in S} x_i = (4l-4)\hat{p}$, since it holds that $\sum_{i \in S} x_i = \hat{p}$. However, since there are l^2 partitions in total, it follows, using a pigeon hole argument that there must be one partition j_0 with $C_{j_0} \leq \frac{4l-4}{l^2} \cdot \hat{p} \leq \frac{4}{l} \cdot \hat{p}$. By building only the base stations of the optimal solution which are not in R_{j_0} we will achieve a profit of at least $\frac{l-4}{l}\hat{p}$. If we consider the new problem instance I_{j_0} in which all base station locations in R_{j_0} are deleted, then the solution of the knapsack problem for partition P_{j_0} is the optimal solution for I_{j_0}. We have already seen that we can solve this knapsack problem with a relative error $\leq 1/l$. So using the partitioning P_{j_0} we find a solution which has an overall profit $\geq \hat{p} \cdot \left(1 - \frac{4}{l}\right)\left(1 - \frac{1}{l}\right) \geq \hat{p} \cdot \left(1 - \frac{4}{l} - \frac{1}{l}\right) = \hat{p} \cdot \left(1 - \frac{5}{l}\right)$.

Hence the profit computed by our algorithm for partition P_{j_0} is an approximation of $\hat{p}$ with relative error $\leq 5/l$. Since we have chosen the largest profit over all partitions, we obtain at least the value $(1 - 5/l)\hat{p}$.

Since l can be chosen beforehand, we see that for every approximation ratio $\epsilon > 1$ (respectively, for every relative error > 0) we obtain a corresponding polynomial time algorithm A_ϵ for the BSPC^C problem. Hence we have proven the following theorem.

THEOREM 1. *For every $C \geq 0$, the problem* BSPC^C *admits a PTAS.*

4. HARDNESS RESULTS

The geometric disk cover problem (DC) is defined as follows: Given is a set N of points in the plane, a radius $r > 0$, and a budget $b \in \mathbb{N}$. The question is if there is a set B of points in the plane, $|B| \leq b$, such that every element from N is covered by a disk with radius r centered in one of the points in B. (If a point j lies in distance exactly d from some $i \in B$, i.e., on a circle with radius r around i, we stipulate that it is covered by a disk centered in i.) It is known that DC is NP-complete [5], even in the strong sense [7].

LEMMA 2. *For $C \geq 12$, the problem* BSPC^C *is strongly* NP-*hard.*

PROOF. First, we show that BSPC^C is NP-hard by constructing a Turing-reduction from DC.

Given is the set N of points in the plane, the radius r, and the budget b. We fix an arbitrary triangulation T which consists of equilateral triangles with side length $\sqrt{3}\cdot r$. (Note that this can be fixed by two points of the plane). It is easy to see that the vertices of T (interpreted as centers of disks) allow a covering of the plane by disks of radius r. Moreover, it can be shown that each circular area with radius r (respectively, $2r$) can be covered by 5 (respectively, 12) disks with centers in T and radius r. For a given circular area with center i and radius r (respectively, $2r$), a minimal covering set $T(i) \subseteq T$ can be determined in polynomial time.

We set $D = r$ and define B' to be the set of all points i in the plane such that if we draw a circle with center i and radius D, then at least two elements from N are located on this circle. Then $|B'| \leq 2\binom{|N|}{2}$. Let B be the union of B' and all $T(i)$ for $i \in B'$. Moreover, let $c(i) = 1$, $e = 1$, $r_{\max}(i) = r$, $o(i) = 0$ and $h(i,k) = 0$ for $i \in B$, $k \in \mathbb{N}$.

Let us consider an optimal solution of this instance. We want to see that in any circular area with radius D there are no more than $C = 12$ built base stations. Otherwise there would exist a circular area with center i and radius r that contains more than 12 base stations. Of course the area that is covered by these base stations is contained in the circular area with center i and radius $2r$. So we would obtain a better solution if we replace these base stations by $T(i)$. This is a contradiction to our assumption and it follows that in the optimal solution in any circular area with radius D there are no more than $C = 12$ built base stations. This shows that for our reduction the values $C \geq 12$ cause no restriction.

Now it is clear that in order to obtain maximal profit, we have to supply as many nodes in N as possible with the given budget. Hence all nodes in N can be covered using at most b disks of radius r iff all nodes in N can be supplied by building b base stations with range D iff the maximal profit is $|N|$.

Besides those numbers already present in the given DC instance, the only non-constant numbers occurring in a constructed BSPC^C instance are the coordinates of the points in B, which can be written down with a number of bits proportional to those needed for coordinates in the input set N, and the value b, which is certainly at most $|B|$, hence bounded by a polynomial in the input length. Therefore, there is a polynomial p such that BSPC^C_p is NP-hard, and thus, BSPC^C is strongly NP-hard. $\square$

THEOREM 3. *For $C \geq 12$, the problem* BSPC^C *does not admit an FPTAS.*

PROOF. The value of an optimal solution of an instance of BSPC^C_p, for p as in the proof of Lemma 2, is at most $|N|$, hence polynomial in the input length. Thus, this problem does not have an FPTAS by [1, Corollary 3.19, p. 116]. $\square$

Hardness of a Restricted Problem

Let us come back to requirement (R3) in Subsect. 1.2. The definition of the problem BSPC^C uses (R3'), which is slightly weaker than what is required in (R3). Therefore, the question remains if the practically more relevant BSPC^{d,D} problem, i.e., restricted to requirement (R3), admits even an FPTAS. In the reduction from disk cover, given in the proof of Lemma 2, we were easily able to conclude that a constant C, as required in the definition of BSPC^C, exists. However, the reduction does not guarantee (R3). Hence, the above reduction does not show that BSPC^{d,D} is hard.

Nevertheless, it can be shown, using a modification of the reduction from 3SAT to DC given in [5], that the restricted BSPC^{d,D} problem is strongly NP-hard; hence we conclude:

THEOREM 4. *For $D \geq d$ the problem* BSPC^{d,D} *does not admit an FPTAS, unless* P = NP.

PROOF. (Sketch.) Fix d, D, where $D \geq d$. Given a 3-CNF formula Φ, we first construct an instance of DC with loops for all variables in Φ as in [5, Theorem 4, pp. 136]. Let N be the set of nodes on these loops. The proof given in that paper shows that all nodes in N can be covered by b disks of suitable fixed radius d iff Φ is satisfiable, where b is the value $N_{\min}$ determined in the proof from [5]. Moreover, a set B of possible disk centers can be determined, given Φ, such that, if Φ is satisfiable, then all nodes can be covered using b disks centered at some of the points in B. A careful analysis shows that any two points in B are at distance at least d.

Now, the reduction from 3SAT to BSPC^{d,D} maps a given formula Φ to the BSPC^{d,D} instance, given by N, B, b as above, and $c(i) = 1$, $e = 1$, $r_{\max}(i) = d$, $o(i) = 0$ and $h(i,k) = 0$ for $i \in B$, $k \in \mathbb{N}$. As in Lemma 2, in order to obtain maximal profit, we have to supply as many nodes in N as possible with the given budget. Hence Φ is satisfiable iff all nodes in N can be supplied by building b base stations with range r iff the maximal profit is $|N|$.

This proves that BSPC^{d,D} is NP-complete. That the problem is actually *strongly* NP-complete follows from the following facts: Coordinates for points in B and N can be described (as in [5]) using small integers, the value of b is at most $|N|$, and all other numbers r, c, e, o, and h are constant.

That $\mathrm{BSPC}^{d,D}$ does not have an FPTAS follows, since even the subset of $\mathrm{BSPC}^{d,D}$, used in the above reduction, is strongly NP-hard, and that the maximal profit is bounded by $|N|$. $\square$

5. CONCLUSION

We examined the complexity of finding an optimal way to build base stations in order to supply a specified demand of teletraffic in a CDMA network. The meaning of "optimal" here took a number of effects in account, such as gain from supplied users, construction costs and ongoing costs for base stations, and higher customer satisfaction if the intra-cell interference is low. We proved that restricted versions of the problem, which are still good enough for practical purposes, have a polynomial-time approximation scheme but no fully polynomial-time approximation scheme.

The sets N_i of demand nodes supplied by a base station i were for our above algorithm just the sets of demand nodes that are located within a certain distance from station i. Certainly this is unrealistic; e.g., it might be that a demand node is located near to a base station, but nevertheless it is not supplied because there is a high building in between causing an area with no reception. However, for our approximation scheme given in Sect. 3, it is only important that (R3') holds. The N_i itself can be *of arbitrary shape*, for example taking skyscrapers or other particularities of the landscape into account. Hence we *do not rely on an Euclidean metric for station signals*. Our algorithm is still polynomial time, if for a given station i and a demand node j it can be determined in polynomial time, if j is supplied by i.

In practice, it is not clear if cells of different base stations can be considered as independent as we did in our model. Even if dependencies between cells have to be taken into account, we suspect that the problems keep their good approximability properties (combining the algorithm of this paper with the method of [6]).

If we look at an *incremental version* of the base station positioning problem, where as an additional input we get a network of already existing base stations that we have to extend by building base stations up to a given budget, it can be observed that a modification of our algorithm shows that this problem also has a PTAS. The only point that changes in the algorithm in Sect. 3 is the computation of the local profits $P(i, m)$. An interesting topic for further research could be to examine a "build on demand" problem, where new teletraffic requests arise online, and we have to update our network continuously.

6. ACKNOWLEDGMENTS

We are very grateful to Phuoc Tran-Gia and Kenji Leibnitz for giving us a tutorial about CDMA technology. During the discussions in these meetings, the formal definition of the BSPC problem evolved. We also acknowledge helpful discussions with Klaus W. Wagner.

7. REFERENCES

[1] 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.

[2] P. Bose, C. Kaklamanis, L. Kirousis, E. Kranakis, D. Krizanic, and D. Peleg. Station layouts in the presence of location constraints. In *Algorithms and Computation, 10th International Symposium*, volume 1741 of *Lecture Notes in Computer Science*, pages 269–278, Berlin Heidelberg, 1999. Springer Verlag.

[3] A. E. F. Clementi, P. Penna, and R. Silvestri. The power range assignment problem in radio networks on the plane. In *Proceedings 17th Symposium on Theoretical Aspects of Computer Science*, volume 1770 of *Lecture Notes in Computer Science*, pages 651–660, Berlin Heidelberg, 2000. Springer Verlag. See also Report No. 54 (2000) of the Electronic Colloqium on Computational Complexity.

[4] S. Eidenbenz, C. Stamm, and P. Widmayer. Positioning guards at fixed height above a terrain – an optimum inapproximability result. In *Algorithms – ESA 98*, volume 1461 of *Lecture Notes in Computer Science*, pages 187–198, Berlin Heidelberg, 1998. Springer Verlag.

[5] R. J. Fowler, M. S. Paterson, and S. L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Information Processing Letters*, 12:133–137, 1981.

[6] C. Glaßer, S. Reith, and H. Vollmer. The complexity of base station positioning in cellular networks. In *Approximation and Randomized Algorithms in Communication Networks, ICALP Workshops 2000*, volume 8 of *Proceedings in Informatics*, pages 167–177. Carleton Scientific, 2000.

[7] 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.

[8] O. H. Ibarra and C. E. Kim. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM*, 22(4):463–468, 1975.

[9] 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.

[10] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman. Analysis of a local search heuristic for facility location problems. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1–10, 1998.

[11] A. Koster, C. van Hoesel, and W. Kolen. Optimal solutions for a frequency assignment problem via tree-decomposition. In *Graph-Theoretic Concepts in*

Computer Science, volume 1665 of *Lecture Notes in Computer Science*, Berlin Heidelberg, 1999. Springer Verlag.

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

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

[14] T. S. Rappaport. *Wireless Communications – Principles and Practice*. Prentice Hall, Upper Saddle River, NJ, 1996.

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

[16] 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.