Planning UMTS Base Station location: optimization models with power control and algorithms*

Edoardo Amaldi, Antonio Capone and Federico Malucelli
DEI, Politecnico di Milano
Piazza L. da Vinci 32, 20133 Milano, Italy
amaldi, capone, malucell@elet.polimi.it

Abstract

Classical coverage models, adopted for second generation cellular systems, are not suited for planning UMTS base station location because they are only based on signal predictions and do not consider the power control mechanism, the traffic distribution and the signal quality constraints. In this paper we propose discrete optimization models and algorithms aimed at supporting the decisions in the process of planning where to locate new base stations. The models consider the Signal-to-Interference Ratio (SIR) as quality measure and capture at different levels of detail the signal quality constraints and the specific power control mechanism of the CDMA air interface. Given that the UMTS base station location problem is NP-hard, we propose two randomized greedy procedures and a Tabu Search algorithm for the uplink (mobile to base station) direction which is the most stringent one from the traffic point of view in the presence of balanced connections such as voice calls. The different models, which take into account installation costs, signal quality and traffic coverage, and the corresponding algorithms are compared on families of small to large-size realistic instances.

1 Introduction

With the extraordinary success of mobile communication services, service providers have been affording huge investments to develop network infrastructures. Due to the high costs and the scarcity of radio resources, an accurate and efficient mobile network planning appears of outmost importance. With the rapid growth of network dimensions and number of users, efficient quantitative

*This work has been supported by "Progetto Cofinanziato MURST 1999, Optimization Models and Methods for Telecommunications Network Design and Management"

methods to support decisions for base station location have become essential. This need is now becoming even more acute with the advent of third generation systems, such as UMTS (Universal Mobile Telecommunication System), due to the increased complexity of the system and the number of parameters that must be considered [7, 21, 3].

The problem of planning second generation cellular systems adopting a TDMA-based access scheme has usually been simplified by subdividing it into a coverage planning problem and a frequency planning problem which are driven by a coverage and, respectively, a capacity criterion [25, 7, 3]. In the coverage planning phase, base stations (BSs) are placed so that the signal strength is high enough in the area to be served [29, 30, 20, 22, 15]. This step only makes use of propagation models, such as for instance Hata's model, to predict the signal levels (see for instance [26]). In the frequency planning phase, a number of channels has to be assigned to each BS [14, 5] so as to take into the traffic load and the level of service measured as the Signal-to-Interference Ratio (SIR).

With the W-CDMA (Wideband Code Division Multiple Access) air interface of UMTS this two-phase approach is not appropriate mainly because the bandwidth is shared by all active connections and no actual frequency assignment is strictly required. The access scheme allows for a more flexible use of the radio resources and the capacity of each cell has not an hard limit (e.g. the number of connections) as in TDMA systems, but it depends on the actual interference levels which determines the achievable SIR values. As these values depend on both traffic distribution and BS positions, BS location in UMTS networks cannot only be based on coverage but it must also be capacity driven [3]. Indeed, interference levels are functions of the emitted powers which, due to a *power control* (PC) mechanism, depend on the user positions. Since the power available at each BS is limited, the number of users that can be served is large if they are close to the BS and small if they are far away (this is the so-called *cell breathing* effect). Therefore, the area actually covered by each BS is heavily affected by the traffic distribution since the number of connections that can be served is limited by the achievable SIR values. It is worth emphasizing that, since interference levels depend both on the connections within a given cell and on those in neighboring cells, the SIR values and the capacity are highly affected by the traffic distribution in the whole area.

The planning phase of cellular networks usually takes as input the following three main types of information related to the service area: i) a set of candidate sites where BSs can be installed, ii) the traffic distribution estimated by using empirical prediction models and iii) the propagation description based on approximate radio channel models or ray tracing techniques. The main purpose of planning is then to select the sites where to install the BSs taking into account different aspects such as costs, signal quality and service coverage.

As we shall see in Section 2, there has been so far little work on optimizing base station location for UMTS networks and, to the best of our knowledge, none of the proposed models considers the impact of traffic distribution, SIR constraints and power control mechanism at an adequate level of detail. In this paper we propose and analyze discrete optimization models aimed at supporting the decisions in the process of planning where to locate the new BSs. Our three models differ in how closely they capture the peculiarities of the power control mechanism of the UMTS CDMA air interface. The focus is on the uplink direction (mobile to base station) which turns out to be the most stringent one from the system capacity point of view in the presence of full-duplex balanced connections such as voice calls. Models for the downlink direction are presented and discussed in [1].

In Section 2 we discuss the main radio network planning issues pertaining to UMTS BS location and comment on previous work. In Section 3 we propose and analyze Integer Linear Programming formulations of this general problem for the W-CDMA setting considering the two most common ways to model the power control mechanism. In Section 4 we describe three heuristics: a randomized greedy and a reverse greedy procedures as well as a Tabu Search algorithm. Computational results obtained with realistic instances are reported and discussed in Section 5. Finally, Section 6 contains some concluding remarks.

2 Radio planning issues for UMTS

UMTS [28] is the third generation mobile communication system developed by ETSI, the European Telecommunications Standard Institute, and is also considered by ITU (International Telecommunication Union) among the standards for the IMT-2000 (International Mobile Telephone standard 2000) family.

One of the two access schemes to be used in the assigned spectrum is based on W-CDMA and frequency division duplexing (FDD). The main characteristic of CDMA is its flexibility in the use of radio resources. In particular, there is no "a priori" limit on the number of simultaneous connections per cell (hard capacity) as with TDMA systems, and resources are dynamically assigned according to interference levels and traffic distribution (soft capacity) [10]. However, this clearly implies an increased complexity in the network planning process and more involved access control procedures. Ad hoc planning and optimization strategies for the CDMA technology are thus needed to actually exploit this additional flexibility [3].

Spreading codes used for signal transmitted in downlink by the same base station are mutually orthogonal, while codes of signals emitted by different stations (base or mobile) can be considered as pseudo-random due to the scrambling sequence [17]. In an ideal environment, the de-spreading process performed at the receiving end can completely avoid the interference of orthogonal signals and reduce that of the others by the *spreading factor* (SF), which is the ratio between the spread signal rate and the user rate. In wireless environments, due to the multipath propagation the interference of orthogonal signals cannot be completely avoided and the SIR is given by:

$$SIR = SF \frac{P_{received}}{\alpha I_{in} + I_{out} + \eta} \tag{1}$$

where $P_{received}$ is the received power of the signal, I_{in} is the total interference due to the signals transmitted by the same base station (intra-cell interference), I_{out} that due to signals of the other base stations (inter-cell interference), α is the orthogonality loss factor ($0 \leq \alpha \leq 1$), and η the thermal noise power. In the uplink case, no orthogonality must be accounted for and $\alpha = 1$.

A simplified model, commonly adopted [32], assumes that the interference due to the neighboring cells (I_{out}) can be expressed as a fraction f of the interference due to the other transmission in the same cell, so that the SIR can be expressed as:

$$SIR = SF \frac{P_{received}}{I_{in}(1+f)} \tag{2}$$

where the thermal noise has not been included assuming that it is much smaller than the interference. This simplified model is usually accurate when the traffic distribution among cells is homogeneous, while it results not appropriate in all the other cases where the contribute of the inter-cell interference can be different for each BS. The values of f usually considered are in the range $0.3 \div 0.5$.

The quality of the received signal, usually expressed in terms of *bit error rate* (BER), mainly depends on the SIR, so it is common to consider quality constraints requiring that the SIR exceeds a minimum value τ which may vary according to the communication service considered (voice, video, packet data, etc.). For the sake of simplicity, in the sequel we refer to the minimum SIR before de-spreading as $SIR_{min} = \tau/SF$.

2.1 Power control and capacity constraints

As mentioned above, the SIR depends on the received powers of the considered signal and the interfering ones. These in turn depend on the transmitted power and the attenuation of the radio link between source and destination. According to propagation conditions, the transmitted power can be adjusted by the power control mechanism so as to minimize interference and guarantee quality. Two PC mechanisms are commonly considered, one based on the received power and one based on the estimated SIR. In the first one, the transmitted power is adjusted so that the power received on each channel is equal to a given target value P_{target}. Similarly in the second one, the transmitted power is set so that the SIR is equal to a target value SIR_{target}. The latter mechanism, adopted for UMTS dedicated channels [28], is more complex since the power emitted by each station depends on that emitted by all the others, but it allows the use of lower powers [33]. Therefore, from a planning prospective, assuming a power-based PC mechanism instead of a SIR-based one leads to a conservative dimensioning which may allocate more radio resources than necessary.

Both in the case of received power-based or SIR-based power control, the transmitted powers are adjusted considering some power limit constraints. In particular, a limit on the maximum power used for each radio channel must me considered both for the uplink and the downlink. Moreover, for the downlink case only, a constraint on the total power emitted by the base station must be added. Therefore, the actual power emitted on a channel is the minimum between that provided by the PC algorithm and the maximum value.

From a planning point of view, the effect of the power bounds and SIR constraints is to limit the capacity of the system. In the power-based PC, as new users are added to the system, the SIR values of all the other users decrease until one falls below the lowest acceptable quality level SIR_{min}. No additional user can then be served. In the SIR-based PC, signal quality is guaranteed by keeping the SIR's at a constant target value $SIR_{target} \geq SIR_{min}$. When new users are added, the emission powers required to keep the SIR's equal to SIR_{target} increase until the power limit would be exceeded and, hence, the SIR would fall below the target value. In both power-based and SIR-based cases, the capacity is affected by user position and propagation conditions. Indeed, capacity depends on the interference generated by user-BS transmissions which in turn depends on the emitted powers which are strongly related to the relative positions and the radio link propagation factors. Since each user-BS transmission generates interference not only in its own cell but also towards all other cells, it is as if each user involved in a connection "absorbs" a fraction of the capacity from all BSs.

2.2 Related work

Classical coverage optimization models do not consider SIR constraints but only constraints on the received power levels in the service area. In [31] it is proposed to describe the traffic distribution by means of *demand nodes* which represent the center of an area characterized by a traffic demand (usually expressed in Erlang). Using the demand node characterization it is possible to define the coverage problem considering the signal level in each node from all base stations and requiring that at least one level is above a fixed threshold. The simplest objective of the optimization process is that of finding the minimum cardinality set of base stations covering all demand nodes (see e.g. [20, 30]). In [23, 24] traffic capacity constraints for each base station are also included to the basic model. A different coverage model is adopted in [29] where the position of transmitters is chosen in a continuous 3D space in order to minimize the sum of path losses of the links to all receivers.

Recently, optimization planning schemes have been also proposed for CDMA systems and for UMTS in particular. However, some works still consider a classical coverage approach: in [4] the minimum dominating set problem is proposed for the modeling of UMTS radio network planning; in [19] the traffic capacity is also considered and the model turns out to be a classical capacitated facility location problem for which a Tabu Search algorithm is proposed. In [6] a different approach is proposed: a graph is introduced where each candidate site is represented by a vertex and where an edge between two vertex exists if the corresponding coverage areas overlap too much. The maximum independent set problem is then used for the base station location.

In [9] a simplified model for locating base stations in CDMA-based UMTS networks which takes into account interference is proposed and a Polynomial Time Approximation Scheme is presented. However, since only the intra-cell interference is considered while the crucial aspect of the interference among base stations is neglected, the interference constraint results in a simple traffic capacity limit.

Since to the best of our knowledge some crucial issues of the planning problem for CDMA-based UMTS networks have not been captured in previous work, we propose and investigate different discrete optimization models that take into account the SIR constraints and the specific aspects of the PC mechanism of CDMA.

3 Base station location models

According to the above observations, the problem of planning an UMTS network appears as a variant of classical *facility location* problems. Besides this common underlying structure, the integer programming models presented here and in [1] differ depending on the transmission direction (uplink or downlink) and the specific type of power control that are considered. In the presence of full-duplex balanced connections (e.g., voice calls), the uplink direction is expected to be the most stringent one from the traffic point of view [10] and, therefore, the most relevant one from a planning point of view. The main focus of this paper is on the location problem with uplink constraints and, unlike in most previous work, our models explicitly account for the PC mechanism (power or SIR based models). Models for the downlink direction are presented in [1].

3.1 Basic model

Consider a territory to be covered by an UMTS service. Assume that a set of *candidate sites* $S = \{1, \dots, m\}$ where a BS can be installed, is given and that an installation cost c_j is associated with each candidate site j, $j \in S$. A set of *test points* (TPs) $I = \{1, \dots, n\}$ is also given. Each TP $i \in I$ can be considered as a centroid where a given amount of traffic d_i (in Erlang) is requested and where a certain level of service (measured in terms of SIR) must be guaranteed [30]. The required number of simultaneously active connections for TP i, denoted by u_i, turns out to be a function of the traffic demand, i.e., $u_i = \phi(d_i)$. The actual definition of the function ϕ is a degree of freedom of the planning process. It can simply correspond to the average number of active connections or to the number of simultaneous connections not exceeded with a given probability p. The connection activity factor can be considered as well.

The propagation information is also supposed to be known. In particular let g_{ij}, $0 < g_{ij} < 1$, be the propagation factor of the radio link between TP i, $1 \leq i \leq n$, and a candidate site j, $1 \leq j \leq m$. The propagation gain matrix $G = \{g_{ij}\}_{1 \leq i \leq n, 1 \leq j \leq m}$ is estimated according to approximate propagation models such as those proposed by Hata or to more precise but computationally intensive ray tracing techniques (see e.g. [26]).

In the W-CDMA UMTS base station location problem one wishes to select a subset of candidate sites within the set S where to install BSs, and to assign the TPs to the available BSs taking into account the traffic demand, the quality requirements in terms of SIR and the installation costs.

Let us define the two following classes of decision variables:

$$y_j = \begin{cases} 1 & \text{if a BS is installed in } j \\ 0 & \text{otherwise} \end{cases}$$

for $j \in S$ and

$$x_{ij} = \begin{cases} 1 & \text{test point } i \text{ is assigned to BS } j \\ 0 & \text{otherwise} \end{cases}$$

for $i \in I$ and $j \in S$. The core of the basic integer linear programming model that we propose for the uplink case is the classical uncapacitated facility location model:

$$\min \quad \sum_{j=1}^{m} c_j \, y_j + \lambda \sum_{i=1}^{n} \sum_{j=1}^{m} u_i \frac{1}{g_{ij}} \, x_{ij} \tag{3}$$

$$s.t. \quad \sum_{j=1}^{m} x_{ij} = 1 \qquad i \in I \tag{4}$$

$$x_{ij} \leq y_j \qquad i \in I; j \in S \tag{5}$$

$$x_{ij}, y_j \in \{0,1\} \qquad i \in I; j \in S. \tag{6}$$

The first term in the objective function corresponds to the total installation cost. Since $1/g_{ij}$ is proportional to the power emitted from TP i when assigned to BS j, the second term aims at favoring assignments which require a smaller total emission power. $\lambda \geq 0$ is a trade-off parameter between these two objectives. Constraints (4) make sure that each TP i is assigned to a single BS. Constraints (5) impose that TPs are only assigned to sites where a BS is installed. Note that by restricting the assignment variables x_{ij} to take binary values, it is required that in every feasible solution all active connections out of any TP must be assigned to a single BS.

To account for the power limit on the user terminals, we need to include, for each pair of TP $i \in I$ and candidate site $j \in S$, the following constraint:

$$\frac{P_{target}}{g_{ij}} \, x_{ij} \leq P_{max} \, y_j \tag{7}$$

where P_{max} is the maximum emission power and P_{target}/g_{ij} corresponds to the emission power required by a mobile station in TP i to guarantee the target received power P_{target} in site j. Note that if $g_{ij}\, P_{max}/P_{target} < 1$ the TP i cannot be assigned to candidate site j due to power limits, and therefore the variable x_{ij} can be omitted from the model. Otherwise constraint (7) is implied by the corresponding constraint (5).

The fundamental aspect to be taken into account is the quality of the signal received by each BS. As mentioned in Section 2, the simplest way to express the quality constraints is either by neglecting the intra-cell interference or by considering that it amounts to a given fraction of the inter-cell interference as given in formula (2) for nonzero values of the parameter f. For each connection, the quality constraint $SIR_{uplink} \geq SIR_{min}$ can then be rewritten as

$$SF\,\frac{P_{received}}{I_{in}(1+f)} \geq SIR_{min} \tag{8}$$

which is equivalent to

$$\frac{I_{in}}{P_{received}} \leq \frac{SF}{SIR_{min}\cdot(1+f)}. \tag{9}$$

Considering a power-based PC mechanism, the power received $P_{received}$ at BS j from each mobile station in a TP assigned to it is equal to P_{target} and the quality constraint amounts to imposing an upper bound on the number N_j of connections that can be assigned to that BS [32]. Specifically, we have

$$N_j \leq \frac{SF}{SIR_{min}\cdot(1+f)}+1. \tag{10}$$

For the typical values $f = 0.4$, $SF = 128$ and $SIR_{target} = 6$ dB used in the literature, we obtain an upper bound of 23.97 on the maximum number of connections that can be serviced by any single BS. Thus for each candidate site $j \in S$ the signal quality constraints can be rewritten as:

$$\sum_{i=1}^{n} x_{ij} \leq 23\; y_j. \tag{11}$$

The resulting model falls within the class of standard capacitated facility location problem which have been extensively studied in the Operations Research literature (see [18]). Medium-size instances of such NP-hard problems turn out to be very difficult to solve even using state-of-the-art optimization techniques. However, the capacity constraints (11) do not capture the distinctive features of the CDMA technology and, as we shall see in Section 5, in most practical cases the above basic model provides meaningless solutions.

3.2 Enhanced model with power-based PC

To make the model more realistic, inter-cell interference needs to be considered explicitly and independently from intra-cell interference. The use of pseudo-random spreading codes implies that, for a specific uplink connection between TP i and BS j, there is no significant difference in the two types of interference.

In the presence of power-based PC mechanism, for each candidate site $j \in S$ we consider the following quality constraint:

$$\frac{P_{target}}{\sum_{h=1}^{n} u_h \, g_{hj} \, \sum_{t=1}^{m} \frac{P_{target}}{g_{ht}} \, x_{ht} - P_{target}} \geq SIR_{min} \, y_j, \tag{12}$$

where the thermal noise η has been omitted as in other works focusing on this type of PC mechanism (see e.g. [32]). This constraint enforces that if a BS is installed in site j (i.e., $y_j = 1$), the resulting SIR must exceed the given SIR_{min}. More specifically, the double summation in the denominator corresponds to the total interference due to all connections and the P_{target} term that is subtracted amounts to the received power associated to a single connection assigned to j. Note that for any TP h the only term of the inner summation (over index t) that is nonzero corresponds to the site to which h is assigned. Thus the contribution to the outer summation of any TP h assigned to site j amounts to $u_h P_{target}$. The bilinear constraints (12) can be linearized as follows:

$$1 + M(1 - y_j) \geq SIR_{min} \cdot \left(\sum_{h=1}^{n} \sum_{t=1}^{m} u_h \, \frac{g_{hj}}{g_{ht}} \, x_{ht} - 1 \right) \tag{13}$$

for a large enough value of M.

It is worth emphasizing that constraints (13) are always satisfied when $y_j = 0$ regardless of the way the TPs are assigned to the BSs, and when $y_j = 1$ can be restated as

$$\sum_{h=1}^{n} \sum_{t=1}^{m} a_{ht}^{j} x_{ht} \leq \frac{1}{SIR_{min}}. \tag{14}$$

Note that $a_{ht}^{j} = u_h - 1$ for one of the TP h assigned to BS j and $a_{ht}^{j} = u_h$ for all other TPs assigned to BS j. For all TPs h assigned to other BSs we have $0 \leq a_{ht}^{j} \leq u_h$. This is clearly in contrast with standard capacity constraints arising in classical capacitated facility location problems that can be expressed as:

$$\sum_{h=1}^{n} d_h x_{hj} \leq s_j y_j \quad j = 1, \ldots, m, \tag{15}$$

where the "demand" d_h of "client" h does not depend on "facility" to which "client" h is assigned (there is no summation on the second subscript of the x variables) and, for each "facility" j, the constraint only involves the "clients" assigned to it [18]. In the context of the UMTS BS location problem, it is as if each TP h absorbs part of the capacity of every BS and not only that of the BS it is assigned to. Moreover, the amount of capacity requested from BS j depends on the "distance" (gain) between TP h and BS j as well as on the "distance" (gain) between TP h and the BS it is assigned to.

By analyzing the quality constraints (12), one can establish a simple but important property of the underlying assignment subproblem of the BS location problem [2]. Given any set of active BSs $\overline{S} \subseteq S$ and set I of TPs that can all be serviced, only assignments of TPs to "closest" active BSs need to be considered, where "closeness" is meant in terms of the required emission power. Indeed, if any one of the TPs is not assigned to one of its "closest" BS the SIR level at each active BS can only remain unchanged or increase when that TP is reassigned to a closer BS. Note that also this property makes the problem under consideration substantially different from standard capacitated facility location problems where, for any given set of active BSs, the optimal assignment is not known a priori and needs to be determined.

3.3 Model with SIR-based PC

Considering a more sophisticated SIR-based PC mechanism, the discrete optimization model is more involved since the signal powers p_i, needed to connect TP $i \in I$ to the candidate site it has assigned to, must be considered as explicit variables. The core model is then as in (3-6) except that in the second term of the objective function $1/g_{ij}$ is replaced by p_i. Constraints (7) are replaced by $0 \leq p_i \leq P_{max}$. Moreover, for each pair of TP $i \in I$ and candidate site $j \in S$, the quality constraint is now:

$$\frac{p_i \, g_{ij}}{\sum_{h=1}^{n} u_h \, g_{hj} \, \sum_{t=1}^{m} p_h \, x_{ht} - p_i g_{ij} + \eta} \geq SIR_{target} \, x_{ij} \tag{16}$$

where $SIR_{target} \geq SIR_{min}$. Note that this third order constraint enforces that, if TP i is assigned to site j (i.e., $x_{ij} = 1$), the transmission powers must be tuned so that the resulting SIR exceeds the given SIR_{target}. Unlike in the power-based PC case, in this context the thermal noise η is usually included to guarantee convergence of the closed-loop PC mechanism used in real systems [33].

3.4 Different planning objectives

Although the generic objective function proposed in (3) takes into account the installation costs and the total emission power, the appropriate choice depends on the specific planning objectives. According to the traffic requirements and distribution, the number of candidate sites and their locations as well as the mobile station maximum power, the signal quality constraints (12) or, respectively, (16) can be infeasible. In real world instances where traffic patterns are based on short-mid term predictions, it is likely that the system will be required to service all traffic. On the other hand, when traffic patterns are based on long term predictions a multi-period network planning strategy can be adopted [27] and in the first stages one has to cope with solutions that do not cover all traffic. In this case, it is reasonable to also aim at maximizing the fraction of traffic that is actually serviced. This can be achieved by relaxing the assignment constraints (4) as follows:

$$\sum_{j=1}^{m} x_{ij} \leq 1 \qquad i \in I \tag{17}$$

so as to allow some TPs not to be assigned. The additional term $\sum_{i \in I} u_i \sum_{j \in S} x_{ij}$ can then be included in the objective function with an appropriate negative weight parameter.

4 Heuristic algorithms

Since the UMTS BS location problems defined in the previous section contain the uncapacitated facility location problem as a special case, they turn out to be NP-hard. To obtain good approximate solutions within a reasonable amount of computing time, we have developed greedy and reverse greedy randomized procedures (see [8]) that iteratively activate and, respectively, disactivate BSs as well as Tabu search heuristic. The focus here is on power-based as well as SIR-based PC mechanisms in the uplink direction.

We first describe the algorithms for models with power-based PC. According to the previous section, we know that in this case assigning each TP to a "closest" active BS yields the largest possible SIR value for each BS. Moreover, in this setting all connections to a given BS j have the same SIR level, denoted by SIR_j. Given any set of active BSs, if all traffic requests can not be covered, the central procedure of our algorithms aims at satisfying the largest fraction of demands (see the relaxed constraints (17)). This subproblem amounts to a special case of multi-knapsack problem which is hard to solve optimally. Since in our context, we need reasonably good solutions in a short amount of time we proceed as follows. Given a set $\overline{S}$ of active BSs, each TP is first assigned to one of its "closest" BS j in $\overline{S}$. Then TPs assigned to active BSs with $SIR_j < SIR_{min}$ are considered by nonincreasing value of emission power $u_i P_{target}/g_{ij}$ and deleted one at a time until all active BSs have a SIR of at least SIR_{min}.

4.1 Randomized greedy procedures

In the first randomized greedy algorithm (Add), BSs are added iteratively to $\overline{S}$ starting from $\overline{S} = \emptyset$. Given the current set $\overline{S} \subseteq S$ of sites in which a BS has already been installed, for each remaining candidate site $j \in S \setminus \overline{S}$ the above procedure is applied to $\overline{S}' = \overline{S} \cup \{j\}$ so as to obtain a corresponding assignment vector $\mathbf{x}$. Each pair $(\overline{S}, \mathbf{x})$ fully characterizes a solution of the original problem since it implies the values of all x and y variables. For each of these solutions $\overline{S}'$, the following utility function is evaluated:

$$U_a(\overline{S}, j) = \sum_{h=1}^{n} \sum_{t=1}^{m} u_h\, x_{ht} - \omega \sum_{t=1}^{m} c_t y_t, \tag{18}$$

where the first term amounts to the fraction of traffic that is currently covered and the second one expresses the total cost for installing the BSs in the sites selected so far. The trade-off parameter ω, $\omega > 0$, allows us to assign higher priority to maximize the first objective than to minimize the second one. At each iteration one $j' \in S \setminus \overline{S}$ is randomly selected among the ρ_a percent ones that yield the largest value of U_a. The procedure stops when the addition of a new BS worsen the current solution value according to utility function U_a.

In the randomized reverse greedy algorithm (Remove), BSs are removed iteratively starting from $\overline{S} = S$. Given the current $\overline{S} \subseteq S$, for each candidate site $j \in \overline{S}$ the above procedure is applied to $\overline{S} \setminus \{j\}$ so as to obtain a corresponding $\mathbf{x}$. Then the following utility function is evaluated:

$$U_r(\overline{S}, j) = \sum_{h=1}^{n} \sum_{t=1}^{m} u_h\, x_{ht} - \omega \sum_{t=1}^{m} c_t y_t + \nu F(\mathbf{x}), \tag{19}$$

where $F(\mathbf{x})$ is the sum, over all BSs j that have so far been installed, of the number of additional connections they could service, namely, of $\lfloor \Delta_j / P_{target} \rfloor$ where Δ_j is defined as the difference between the current SIR and SIR_{min}. $\nu > 0$ is the corresponding weight. At each iteration one $j' \in \overline{S}$ is randomly selected among the ρ_r percent ones that yield the largest value of U_r. As for the

Add procedure, the Remove procedure stops when the removal of another BS worsen the current solution value according to utility function U_r.

Given the randomized nature of the Add and Remove procedures and their relatively low computational requirements, a multistart strategy is adopted: they are run a predefined number of times and provide as output the best solution found.

4.2 Tabu Search

Tabu Search (TS) is a *metaheuristic* that guides a local search procedure to explore the solution space of optimization problems beyond local optima. The idea is to use the history of the search process through an appropriate memory scheme to prevent cycling (running into feasible solutions that have already been generated) and to explore regions of the solution space that are promising in terms of the objective function. The modern TS paradigm goes back to the seminal work by Glover [11, 12] and is extensively discussed in [13].

The basic ingredients of a general TS strategy can be described as follows. Starting from an initial feasible solution $\mathbf{s}^0$, a set of neighboring solutions $N(\mathbf{s}^0)$ are generated by applying a set of possible "moves" to $\mathbf{s}^0$. Then the best solution in the "neighborhood" $N(\mathbf{s}^0)$ is selected as the next iterate $\mathbf{s}^1$ even if it does not strictly improve the value of the objective function and the process is repeated to generate a sequence of solutions $\{\mathbf{s}^k\}$.

In order to prevent cycling and to try to escape from local optima a list of "tabu moves" is maintained. The purpose of this list is to prohibit to make the opposite move that has been made at a given step for a certain number of iterations. A move that is added to the list remains tabu for a number of iterations that corresponds to the length L of the list. According to the "aspiration criteria", tabu moves can be clearly made if they lead to an improved solution. The best solution encountered is stored as the algorithm proceeds and it is returned after a maximum number of iterations max_{it}. max_{it} and L are two parameters of the method. Not surprisingly the efficiency and performance of a TS procedure strongly depend on the way the moves are defined and how well they exploit the actual structure of the problem at hand.

As initial solution we consider the set of active BSs $\overline{S} \subseteq S$ provided by the Add or the Remove procedures together with an assignment $\mathbf{x}$ of the TPs to their closest active BS, whenever they are serviced. Note that the current solution is in a sense fully characterized by the set $\overline{S} \subseteq S$ of active BSs. Besides the *add* and *remove* moves used in the randomized greedy procedures, we do also define *swap* moves that amount to installing a new BS in one of the empty sites while deleting one of the active BSs. Exploring all possible swap moves for any given current solution would of course be very time consuming even for small-size instances. But, due to the actual structure of the problem, it is reasonable to focus on swaps between candidate sites that are relatively "close" to each other. For each active BS j we initially consider the set $C \subseteq S$ of max_{swap} available sites that have the largest propagation gains with respect to j and that are in a sense the best candidates for a swap. Let q be a parameter such that $0 < q < 1$. The swaps involving the $max_{swap} \cdot q$ best sites in C are systematically evaluated, while those corresponding to the remaining available sites in C are considered with a given probability q'.

The objective function used to guide our Tabu Search is U_r given in (19). As to the implementation of the tabu list, BSs that are installed (disactivated) cannot be disactivated (reinstalled) during L iterations. Although this imposes stronger restrictions on the search process than just avoiding considering solutions that have already been generated, it provides better experimental results since it forces to explore a larger region of the solution space (see the concept of diversification [13]).

In the experiments described in Section 5, Tabu Search is applied in two different settings. On the one hand, we use 200 TS iterations in a multistart Add or Remove context as a local search procedure to improve the solutions obtained with each one of the 10 runs of the greedy method. On the other hand, the same total number of iterations, namely 2000, of TS are carried out starting from a single initial solution provided by Add or Remove. Note that the advantage of Remove to generate a starting solution is that it allows for automatically checking whether all traffic demand can be satisfied.

4.3 Extension for the model with SIR-based PC

Although in the SIR-based case assigning TPs to closest BSs is no longer guaranteed to yield the best possible SIR values, we make this simplifying assumption which reflects a conservative point of view in the sense that there might exist an assignment of TPs to active BSs involving a smaller number of active BSs. Given a current set of active BSs $\overline{S} \subseteq S$ and an assignment $\mathbf{x}$, to compute the emitted powers p_i it would be necessary to solve the system composed of constraints (16) for each active connection between TP and BS together subject to the bounds on the maximum power.

Since the resulting system contains n constraints (16) in the n variables p_i, a solution satisfies all of them with equality. Assuming that TPs are assigned to closest active BSs, the system can then be simplified by reducing the number of variables as well as the number of equations. For every TP i assigned to a given active BS j, all corresponding equations are equivalent when considered as function of $p_i g_{ij}$. Denoting the common received power by $p'_j = p_i g_{ij}$, the equations can be rewritten as:

$$\frac{p'_j}{\sum_{h=1}^{n} u_h\, g_{hj}\, \sum_{t=1}^{m} \frac{p'_t}{g_{ht}}\, x_{ht} - p'_j + \eta} = SIR_{target} \quad \text{for all } j \in \overline{S}. \tag{20}$$

The size of the resulting system is thus equal to the number of active BSs which is at most m and usually much smaller than the number n of TPs. This size reduction is absolutely necessary to make our algorithms applicable to the more accurate model. Given the assignment $\mathbf{x}$ and the solution $\mathbf{p}'$ of the the above system, the emission power at the TPs are derived by setting $p_i = p'_j/g_{ij}$ if $p_i \leq P_{max}$ and otherwise $p_i = P_{max}$.

5 Computational results

To evaluate the performance of the proposed algorithms we have considered synthetic but realistic uplink instances generated by using the Hata's propagation model [16]. For each instance we consider a rectangular service area with dimensions $L \times W$, a number m of candidate sites in which to locate omnidirectional antennas and a number n of TPs. Using a pseudo-random number generator each candidate site j and each TP i is assigned a position with uniform distribution in the service area. The matrix $\mathbf{G}$ is obtained by using the Hata's formulas which give the attenuation A (loss) in dB due to signal propagation. In particular, the attenuation for urban areas is given by:

$$\begin{aligned} A_u \quad = \quad & 69.55 + 26.16\log(F) - 13.82\log(H_b) + \\ & - [(1.1\log(F) - 0.7)H_m - (1.56\log(f) - 0.8)] + \\ & + [44.9 - 6.55\log(H_b)]\log d, \end{aligned} \tag{21}$$

where F is the signal frequency in MHz, H_b and H_m are the heights of the base and the mobile station in meters, and d is the distance in Km, while the formula for rural areas is:

$$A_r = A_u - 4.78[\log(F)]^2 + 18.33\log(F) - 35.94. \tag{22}$$

Clearly $g_{ij} = 1/10^{A(d_{ij})/10}$, where d_{ij} denotes the distance between TP i and candidate site j.

We considered three families of uplink instances with a service area of 400×400 m, 1×1 Km and 1.5×1.5 Km respectively. Different instances of these families are obtained using the pseudo-random generator. Two gain matrices $\mathbf{G}$ are considered for each configuration using the urban and the rural Hata's formulas with:

$$H_b = 10\ \ m, \qquad H_m = 1\ \ m, \qquad F = 2000\ \ MHz. \tag{23}$$

Small-size instances are characterized by a service area of 400×400 meters, $m = 22$ candidate sites, $n = 95$ test points, and $u_i = 1$ for all TPs $i \in I$. Medium-size instances are characterized by a service area of 1×1 km, $m = 120$, $n = 400$ and, since u_i uniformly distributed in $\{1, 2, 3\}$ for all $i \in I$, the total number of connections is equal to 800. Large-size urban instances are characterized by a service area of 1.5×1.5 Km, $m = 200$, $n = 750$, and u_i uniformly distributed in $\{1, 2\}$ for all $i \in I$ so that the total number of connections is equal to 1125.

Finally, the remaining parameters have selected as follows: installation costs $c_j = c$ for all $j \in S$, $SF = 128$ and $\tau = 4$ (so $SIR_{min} = 0.03125$), $SIR_{target} = 6$ dB, $P_{max} = 30$ dBm, $P_{target} = -100$ dBm. In all the experiments reported here $\omega = 1/(c\ m)$. and $\nu = 1/m^2$.

5.1 Shortcomings of the simplified SIR formula

The first computational effort has been directed towards a comparison between the simplified power-based PC model (i.e., including constraints (11)) and the enhanced power-based PC one with SIR constraints (12) in which inter-cell interference is explicitly considered. Setting $f = 0.4$ and the other parameters as mentioned above, in the simplified model the maximum number of users per BS is bounded above by 23.97 and, hence, it is at most 23. The resulting models for the small instances have been solved to optimality by using the Mixed Integer Programming solver CPLEX 7.0. In all cases, the simplified model activates 5 BSs (which corresponds to the lower bound, since 95/23 is about 4.1) while the optimal solution of the enhanced model always activates just 4 BSs.

To evaluate the quality of the solutions yielded by the simplified model, we consider the active BSs and the assignments of TPs provided by CPLEX, and we compute the SIR value of each BS using a power-based PC mechanism as well as the power received by each BS using a SIR-based PC mechanism. Typical results obtained for one of the small instances of Table 3 (instance SU-3) are reported in Table 1. SIR values and received powers are given for each of the five active BSs. Obviously in the power-based PC setting the power received is constant, while in the SIR-based setting the SIR is constant. In the power-based setting SIR values often exceed the minimum level required to guarantee signal quality, while in the SIR-based setting the SIR values are set to $SIR_{target} = 6$ and the actual received powers associated to a number of connections can be lower than P_{target}.

BS id.	**TPs**	**power-based setting** **SIR (dB)**	**Power (dBm)**	**SIR-based setting** **SIR (dB)**	**Power (dBm)**
5	18	6.21	-130	6	-134.26
14	13	6.42	-130	6	-134.79
15	21	6.77	-130	6	-135.47
19	22	6.21	-130	6	-134.34
20	21	6.46	-130	6	-134.73

Table 1: Quality of the solutions from the simplified model with $f = 0.4$ for instance SU-3 of Table 3.

For a fair comparison between the solutions provided by the simplified and enhanced models, we changed the value of parameter f (setting it to 0.35) so that also the simplified model returns solutions with only 4 active BSs. The corresponding results obtained for the same instance of Table 1 are reported in Table 2. Notice that this time in the power-based PC setting the SIR values drop far below any acceptable value, while in the SIR-based PC setting it has not even been possible to find a feasible value for the received powers, meaning that the PC mechanism does not converge to a stable solution.

BS id.	**TPs**	**power-based setting** **SIR (dB)**	**Power (dBm)**
0	23	4.8	-130
6	24	4.97	-130
8	24	5.09	-130
13	24	5.22	-130

Table 2: Quality of the solutions of the simplified model with $f = 0.35$ for instance SU-3.

Interesting insight on the inadequacy of the location model with the simplified SIR constraints is given by the distribution of the number of connections assigned to each active BS in the solutions provided by the other more sophisticated models. A typical example for a medium-size instance (instance MU-3) is shown in Figure 1. In the case of the SIR-based PC model, a substantial fraction of BSs (here more than 37%) is assigned a number of connections that is larger than 23, which is the strict upper bound that would be imposed by the simplified power-based PC model.

Solutions of the enhanced power-based PC model do also exhibit similar distributions.

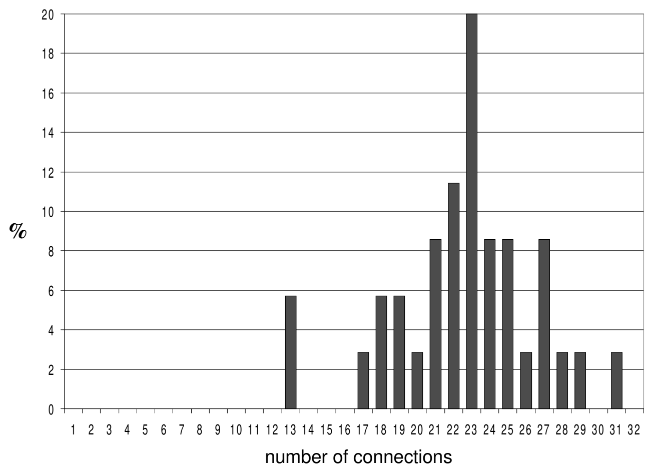

Figure 1: Distribution of the number of TPs assigned to the various BSs for the medium-size instance MU-3.

All our experiments show that the model with simplified SIR formula is not suitable for locating BSs in UMTS networks since it fails to capture the fundamental characteristics of third generation systems.

5.2 Comparative results for model with power-based PC

In this subsection the results obtained with the three algorithms for the enhanced power-based PC model with thermal noise $\eta = 0$ are reported and discussed.

Results for the small-size instances obtained with the Add and Remove randomize procedures ($\rho_a = \rho_r = 30\%$) over 50 runs are reported in Table 3. The solutions found always satisfy all 95 requested connections. Out of the 10 instances mentioned in Table 3, 5 have been generated using the Hata formula for urban areas (SU) and 5 using that for rural areas (SR). For the same instances the optimal solutions, involving 4 BSs in all cases, have been obtained by using CPLEX 7.0. In half of the cases, Add and Remove provide an optimal solution and the number of BSs installed never exceeds the minimum number by more than 1. It is worth noting that the computation time required by Add or Remove for a single instance was less than 5 seconds on a Pentium III/700 MHz Personnal Computer, while CPLEX required 5 to 20 minutes on an about twice as fast computer. As a matter of fact, solving even such small instances using a state-of-the-art MIP solver turns out to be a challenging task due to the very delicate choice of the value M. If the parameter M is too small, then some SIR constraints that should be omitted will not be disactivated and, therefore, the resulting solutions contain a larger number of BSs than optimal solutions of the actual problem. On the other hand, for too large values of M, one runs into numerical problems related to machine precision which tend to override some SIR constraints and, hence, to yield solutions with a too small number of BSs.

	Add			Remove			multistart TS (A)			multistart TS (R)			TS
SU-1	5	5.80	0.85	4	5.22	0.46	4	4.62	0.49	4	4.42	0.49	4
SU-2	5	5.22	0.41	4	4.74	0.44	4	4	0	4	4	0	4
SU-3	4	5.32	0.58	4	5.08	0.74	4	4.02	0.14	4	4	0	4
SU-4	5	6.72	0.96	5	5.58	0.64	5	5	0	5	5	0	5
SU-5	4	5.06	0.54	4	4.90	0.30	4	4	0	4	4	0	4
SR-1	5	5.74	0.69	4	5.34	0.65	4	4.50	0.50	4	4.44	0.49	4
SR-2	4	5.16	0.50	4	4.72	0.45	4	4	0	4	4	0	4
SR-3	4	5.18	0.51	4	4.72	0.49	4	4.02	0.14	4	4	0	4
SR-4	5	5.60	0.90	5	5.34	0.62	5	5	0	5	5	0	5
SR-5	4	5.04	0.45	4	5.02	0.37	4	4	0	4	4	0	4

Table 3: Results obtained with Add, Remove ($\rho_a = \rho_r = 30\%$), multistart TS and single run TS for small-size urban (SU) and rural (SR) instances with $m = 22$, $n = 95$, $u_i = 1$. For each algorithm, from left to right: minimum number of BSs installed, the average number and the standard deviation.

In Table 3, multistart TS consists of 50 runs of Add (A) or Remove (R) followed by 200 iterations of Tabu Search while single run TS consists of Remove followed by 2000 iterations of Tabu Search. For these small instances, a tabu list of length $L = 8$ is used, $max_{swap} = 5$, $q = 1$ and $q' = 0$. As far as the best solutions are concerned, for these small instances the Remove procedure provides as good solutions as the TS variants. However, in multistart TS and single run TS the average quality of the solutions is better and the standard deviations are lower in all but one case. Due to this more robust behavior, a single run of multistart TS would suffice to yield the same best solutions for 8 out of the 10 instances.

Table 4 reports the results obtained with Add, Remove (both with 50 runs) and multistart TS (Add and Remove) and single run TS for 10 medium-size instances. The number of TS iterations for multistart and single run TS are as for small-size instances. For medium-size instances, a tabu list of length $L = 15$ is used, $max_{swap} = 15$ and $q = q' = 0.3$. The solutions found always provided full traffic coverage. The number of BSs installed by our Add and Remove randomized procedures differ at most by 3. None of them consistently outperforms the other one. However, the standard deviation of the number of BSs in the solutions given by Remove (over the 50 runs) is usually higher than that of the solutions provided by Add. The solutions provided by the TS procedures require a few BSs less than Add and Remove. Although there is no significant difference among the three TS procedures, single run TS provides in all but one cases the best solution and it always does so at lower computational cost.

	Add			Remove			multi TS (A)			multi TS (R)			TS
MU-1	47*	50.30	1.53	50	56.16	7.36	46	46.90	1.04	48	48.50	0.50	47
MU-2	46	48.26	1.51	46	49.06	6.65	43	43.40	0.66	43	43.40	0.49	43
MU-3	45	46.40	1.18	43	48.98	4.06	41	41.00	0	41	41	0	41
MU-4	45	48.10	1.36	44	48.08	2.98	42	42.30	0.46	42	42.10	0.3	42
MU-5	44	47.54	1.37	46	49.86	3.26	42	42.30	0.46	42	42.10	0.3	42
MR-1	44	45.88	1.13	42	45.80	3.23	40	40.80	0.40	41	41	0	40
MR-2	44	47.72	1.48	45	48.60	5.72	43	43.10	0.30	43	43.30	0.46	43
MR-3	43	46.26	1.60	44	48.80	4.83	41	41.10	0.30	41	41.20	0.4	41
MR-4	45	48.34	1.57	45	49.08	4.5	42	42.10	0.30	42	42	0	42
MR-5	44	46.38	1.13	46	49.64	2.64	42	42.40	0.49	42	42.30	0.46	42

Table 4: Results obtained with the Add, Remove ($\rho_a = \rho_r = 0.3$), multistart TS and single run TS for medium-size urban (MU) and rural (MR) instances with $m = 120$, $n = 400$ and u_i uniformly distributed in $\{1, 2, 3\}$. For each algorithm, from left to right: minimum number of BSs installed, the average number and the standard deviation over 50 runs. * Not all the traffic is covered.

The 50 runs take in average 10 seconds for Add and 30 seconds for Remove for the small-size instances, 1:50 and, respectively, 1:20 for the medium-size instances and for large-size instances more than 13 hours for Add and approximately 7 hours for Remove. A single run TS lasts in average 1:20 hours for medium-size instances and 6 hours for large-size ones. For solution quality comparison purposes, multistart TS was run with the same total number of Tabu Search iterations as single run TS and obviously this led to higher (approximately twice as long) computing times.

As shown in Table 4, the multistart TS (10 runs) and the single run of TS yield for all instances a solution better than the best ones obtained with Add and Remove. In many cases these solutions are obtained within less than 500 local search moves.

We have also applied Add and Remove (10 runs), multistart TS (10 runs with 200 iterations of local search each) as well as single run TS (1000 iterations) on five large-size urban instances. This choice of parameter values, with 1500 active connections in the average, is quite realistic for UMTS setting in medium-to-large cities. The results are reported in Tables 5 and 6. Unlike for the previous classes of instances, the traffic demand is not always satisfied. It is worth noting that Remove provides solutions with much higher average number of active BSs even though the best solution found is worse than the best one yield by Add only for two out of the five instances. In fact, the best solutions found with Remove turn out to be better for instances LU-3 and LU-5 and equivalent for LU-1. The advantage of the Remove scheme to obtain the initial solutions is that it naturally allows for testing whether all traffic demand can be satisfied.

	Add				Remove			
LU-1	1106/1107	68	70.80	1.89	1106/1107	68	100.34	17.41
LU-2	1133/1136	72	70.46	2.00	1136/1136	73	92.54	9.14
LU-3	1120/1124	64	67.46	2.15	1120/1124	61	73.14	13.10
LU-4	1130/1130	63	66.60	2.04	1130/1130	69	89.08	7.02
LU-5	1103/1103	61	64.66	2.02	1103/1103	60	64.40	2.68

Table 5: Results obtained with Add and Remove ($\rho_a = \rho_r = 0.3$) for large urban instances with $m = 200$, $n = 750$, u_i uniformly distributed in $\{1, 2\}$. From left to right: fraction of traffic covered, minimum number of BSs installed, average number and standard deviation over 10 runs.

	multi TS (A)				multi TS (R)				TS
LU-1	1106/1107	62	63	0.89	1106/1107	64	66.9	1.64	62
LU-2	1136/1136	63	61.3	1.10	1136/1136	63	65.1	1.22	63
LU-3	1120/1124	57	57.5	0.67	1120/1124	57	58.1	1.14	57
LU-4	1130/1130	57	58	0.63	1130/1130	58	61.8	1.6	60
LU-5	1103/1103	54	54.7	0.78	1103/1103	54	54.7	1.01	56

Table 6: Results for large urban instances with $m = 200$, $n = 750$, u_i uniformly distributed in $\{1, 2\}$. From left to right: fraction of traffic covered, minimum number of BSs installed, average number and standard deviation.

According to Table 6, multistart TS and single run TS always yield solutions of better quality than the best ones provided by Add and Remove. For these large-size instances, a tabu list of length $L = 15$ is used, $max_{swap} = 15$, $q = 0.26$ and $q' = 0.3$. Note that for four out of the five instances the improvement in terms of BSs installed is substantial.

5.3 Comparison between models with power-based or SIR-based PC

We now discuss the results provided by the model with SIR-based PC pointing out the differences with respect to those from the models with power-based PC. To obtain meaningful solutions with the SIR-based model a nonzero thermal noise η must be considered. For comparison purpose, the same value of $\eta = -130$ dB is also used in the model with power-based PC.

Multistart Tabu Search is applied with 10 runs. To maintain the computational times at a reasonable level, TPs have been assigned to closest BSs in terms of emitted powers, even though better performance (and possibly also lower costs) may be attained with different types of assignments. At each iteration, once the current solution is fully determined (i.e., a TP-BS assignment is derived for the current set of active BSs) the received powers are computed by solving the related linear system (20). This constitutes the heavy part of the computation. Tests have been conducted on the small-size instances and both algorithms provided solutions with 4 BSs, as in the enhanced power-based PC case. In order to emphasize the differences, experiments were also carried out with the medium-size instances. For each of the medium-size instances the multistart TS took about 15 hours while the single run TS took about 9 hours. These computing times have to be compared to the 2:30 and 1:30 hours respectively for the same type of algorithms applied to the enhanced power-based PC. This clearly indicate the considerable additional computational load required to compute the power values at each iteration. Table 7 shows the numerical results for the ten medium-size instances. For multistart TS the minimum number of installed BSs is reported together with the average and the standard deviation (over the 10 iterations), while for single run TS the best solution is given.

	SIR-based PC				**power-based PC**		
	multistart TS			TS	multistart TS		
Instance	best	average	st. dev.	best	best	average	st. dev.
MU-1	39	39	0	39	52	53.1	0.83
MU-2	37	37.1	0.3	36	45	45.5	0.5
MU-3	35	35.1	0.3	35	43	43.4	0.66
MU-4	36	36.7	0.46	36	43	44.1	0.94
MU-5	36	36.5	0.5	36	43	44.5	1.20
MR-1	35	35.9	0.3	36	42	42.5	0.5
MR-2	37	37	0	36	45	45.5	0.5
MR-3	35	35.1	0.3	35	43	43.3	0.46
MR-4	36	36.4	0.49	36	43	44.7	1.19
MR-5	36	36.6	0.49	36	43	43.9	0.83

Table 7: Comparison between models with SIR-based and power-based PC: number of installed BSs with Tabu Search algorithms

Comparing the number of active BSs yielded by the enhanced power-based PC model with that obtained by the SIR-based PC model we observe that the latter one allows for a saving of at least 6 BSs. This is due to the fact, since the mobile stations can emit lower powers with respect to those in the power-based PC case, they generate lower interference towards all BSs they are not assigned to. Thus the system can support more connections per BSs and, in practice, the capacity of the overall network is better exploited.

The above observations are also confirmed by comparing the SIR values at each BS and the received powers for each connection. Figure 2 reports the SIR values measured in each active BS by using the enhanced power-based PC model. Notice that only a few BSs have a SIR equal to the required SIR_{target} of 6 db, and for many BSs the service quality is far above the required level. Thus the emitted powers are higher than needed and as a result the interference levels are higher and the scarce radio resources are not used as efficiently as they could.

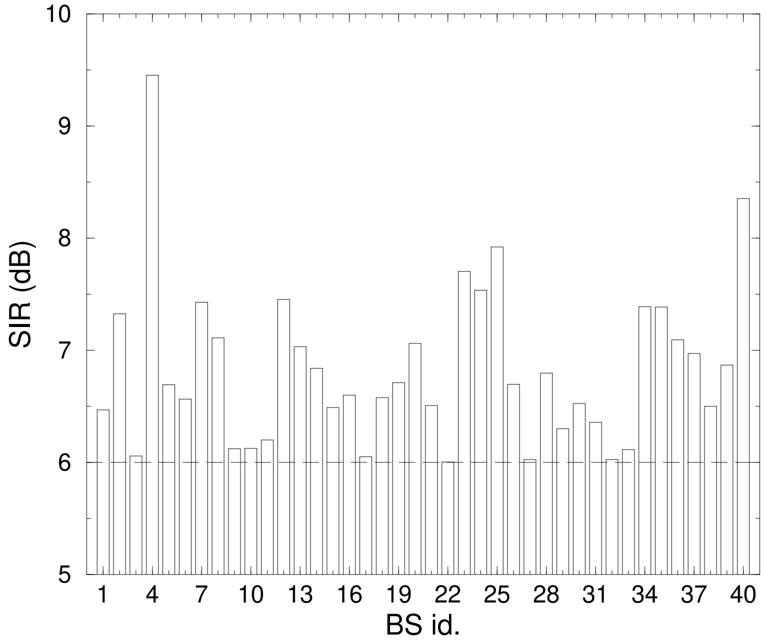

Figure 2: Measured SIR values with the enhanced power-based PC model for instance MU-3.

6 Concluding remarks

The main features characterizing the important problem of planning UMTS networks using the W-CDMA air interface have been discussed. Focusing on the uplink direction, we have investigated discrete optimization models for the UMTS BS location problem that capture at different levels of detail the peculiarities of the signal quality constraints and the PC mechanism for CDMA. Due to the inadequacy of standard location models, we have proposed an enhanced power-based PC model as well as a more accurate SIR-based PC one.

Since solving even medium-size instances of these NP-hard problems is beyond the reach of state-of-the-art commercial optimization packages, we have developed three heuristics for the power-based PC model. Our randomized greedy and reverse greedy procedures provide, in a reasonable amount of time, good approximate solutions for medium-to-large size realistic instances generated with classical propagation models. The feasible solutions obtained with these greedy procedures can then be further improved by applying a Tabu Search method either in a multistart or single run setting.

The three above algorithms have been extended to the SIR-based model by assuming that TPs are assigned to a closest active BS. From the planning point of view this is just a conservative assumption in the sense that there might exist better solutions with less natural assignments but at lower cost.

Our experimental results show that the enhanced power-based PC model yields interesting solutions but those obtained with the SIR-based model use in a more efficient way the scarce radio resources and the computed SIR values are closer to the actual levels of a real system.

Acknowledgments

The authors would like to thank Marco Fumagalli and Stefano Pezzoli for their significant help in the long and tedious experimental phase.

References

[1] E. Amaldi, A. Capone, and F. Malucelli. Base station configuration and location problems in UMTS networks. In *Proceedings of the 9th International Conference on Telecommunication Systems, Modeling and Analysis*, pages 341–348, 2001.

[2] E. Amaldi, A. Capone, and F. Malucelli. Discrete models and algorithms for the capacitated location problem arising in UMTS network planning. In *Proceedings of the 5th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, pages 1–8. ACM, 2001.

[3] E. Berruto, M. Gudmundson, R. Menolascino, W. Mohr, and M. Pizarroso. Research activities on UMTS radio interface, network architectures, and planning. *IEEE Communications Magazine*, pages 82–95, February 1998.

[4] P. Calégari, F. Guidec, P. Kuonen, and D. Wagner. Genetic approach to radio network oprimization for mobile systems. In *Proceedings of IEEE VTC 97*, pages 755–759, 1997.

[5] A. Capone and M. Trubian. Channel assignment problem in cellular systems: a new model and tabu search algorithm. *IEEE Trans. on Vehicular Technology*, 48:1252–1260, 1999.

[6] B. Chamaret, S. Josselin, P. Kuonen, M. Pizarroso, B. Salas-Manzanedo, S. Ubeda, and D. Wagner. Radio network optimization with maximum independent set search. In *Proceedings of IEEE VTC 97*, pages 770–774, 1997.

[7] J.C.S. Cheung, M.A. Beach, and J.P. McGeehan. Network planning for third-generation mobile radio systems. *IEEE Communications Magazine*, pages 54–59, November 1994.

[8] T. A. Feo and M. G. C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6(2):109–133, 1995.

[9] M. Galota, C. Glasser, S. Reith, and H. Vollmer. A polynomial-time approximation scheme for base station positioning in UMTS networks. In *Proceedings of the 5th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (DIAL-M)*, pages 52–59. ACM, 2001.

[10] K. S. Gilhousen, I. M. Jacobs, R. Padovani A. J. Viterbi, L. A. Weaver, and C. E. Wheatley. On the capacity of a cellular CDMA system. *IEEE Trans. on Vehicular Technology*, 40:303–312, 1991.

[11] F. Glover. Tabu search –part i. *ORSA Journal on Computing*, 1(3):190–206, 1989.

[12] F. Glover. Tabu search –part ii. *ORSA Journal on Computing*, 2(1):4–32, 1990.

[13] F. Glover and M. Laguna. *Tabu search.* Kluwer Academic Publishers, 1997.

[14] W. K. Hale. Frequency: Theory and applications. *Proceedings of IEEE*, 30:1497–1514, 1980.

[15] Q. Hao, B.-H. Soong, E. Gunawan, J.-T. Ong, C.-B. Soh, and Z. Li. A low-cost cellular mobile communication system: a hierarchical optimization network resource planning approach. *IEEE Journal on Selected Areas in Communications*, 15:1315–1326, 1997.

[16] M. Hata. Empirical formula for propagation loss in land mobile radio service. *IEEE Trans. on Vehicular Technology*, 29:317–325, 1980.

[17] H. Holma and A. Toskala. *WCDMA for UMTS.* John Wiley and sons, 2000.

[18] M. Labbé and F. V. Louveaux. Location problems. In M. Dell'Amico, F. Maffioli, and S. Martello, editors, *Annotated Bibliographies in Combinatorial Optimization*, chapter 16. John Wiley, 1997.

[19] C.Y. Lee and H.G. Kang. Cell planning with capacity expansion in mobile communications: a tabu search approach. *IEEE Trans. on Vehicular Technology*, 49(5):1678–1690, 2000.

[20] R. Mathar and T. Niessen. Optimum positioning of base stations for cellular radio networks. *Wireless Networks*, pages 421–428, 2000.

[21] R. Menolascino, M. Pizarroso, C. Lepschy, and B. Salas. Third generation mobile systems planning issues. In *Proceedings of IEEE VTC 98*, pages 830–834, 1998.

[22] A. Molina, G.E. Athanasiadou, and A.R. Nix. The authomatic location of base stations for optimized cellular coverage: a new combinatorial approach. In *Proceedings of IEEE VTC 99*, pages 606–610, 1999.

[23] A. Molina, G.E. Athanasiadou, and A.R. Nix. Cellular network capacity planning using the combinatorial algorithm for totatl optimization. In *Proceedings of IEEE VTC 2000*, pages 2512–2516, 2000.

[24] A. Molina, G.E. Athanasiadou, and A.R. Nix. Optimized base station location algorithm for next generation microcellular networks. *Electronic Letters*, 36(7):668–669, 2000.

[25] M. Naghshineh and I. Katzela. Channel assignment schemes for cellular mobile telecommunications systems: A comprehensive survey. *IEEE Personal Communications*, 3 (3):10–31, 1996.

[26] D. Parsons. *The mobile radio propagation channel.* Wiley and Sons, West Sussex, 1996.

[27] P. Reininger, S. Iksal, A. Caminada, and J. Korczak. Multi-stage optimization for mobile radio network planning. In *Proceedings of the IEEE VTC 99*, pages 2034–2038, 1999.

[28] A. Samukic. UMTS universal mobile telecommunications system: development of standards for the third generation. *IEEE Transactions on Vehicular Technology*, 47(4):1099–1104, 1998.

[29] H.D. Sherali, M. Pendyala, and T. Rappaport. Optimal location of transmitters for micro-cellular radio communication system design. *IEEE Journal on Selected Areas in Communications*, 14(4):662–673, 1996.

[30] K. Tutschku. Demand-based radio network planning of cellular mobile communication systems. In *Proceedings of the IEEE INFOCOM'98*, volume 3, pages 1054–1061, 1998.

[31] K. Tutschku, N. Gerlich, and P. Tran-Gia. An integrated approach to cellular network planning. In *Proceedings of the 7th International Network Planning Symposium - Networks 96*, 1996.

[32] A. M. Viterbi and A. J. Viterbi. Erlang capacity of a power controlled cdma system. *IEEE Journal on Selected Areas in Communications*, 11(6):892–900, 1993.

[33] R. D. Yates. A framework for uplink power control in cellular radio systems. *IEEE Journal on Selected Areas in Communications*, 13:1341–1347, 1995.