J Heuristics (2007) 13: 243–263
DOI 10.1007/s10732-007-9015-5

Automated antenna positioning algorithms for wireless fixed-access networks

Yehuda Ben-Shimol · Boaz Ben-Moshe · Yoav Ben-Yehezkel · Amit Dvir · Michael Segal

Received: 16 November 2005 / Revised: 9 April 2006 /
Accepted: 31 July 2006 / Published online: 13 April 2007
© Springer Science+Business Media, LLC 2007

Abstract This article addresses a real-life problem - obtaining communication links between multiple base station sites, by positioning a minimal set of fixed-access relay antenna sites on a given terrain. Reducing the number of relay antenna sites is considered critical due to substantial installation and maintenance costs. Despite the significant cost saved by eliminating even a single antenna site, an inefficient manual approach is employed due to the computational complexity of the problem. From the theoretical point of view we show that this problem is not only NP hard, but also does not have a constant approximation. In this paper we suggest several alternative automated heuristics, relying on terrain preprocessing to find educated potential points for positioning relay stations. A large-scale computer-based experiment consisting of approximately 7,000 different scenarios was conducted. The quality of alternative solutions was compared by isolating and displaying factors that were found to affect the standard deviation of the solutions supplied by the tested heuristics. The results of the simulation based experiments show that the saving potential increases when more base stations are needed to be interconnected. The designs of a human expert were compared to the automatically generated solutions for a small subset of the experiment scenarios. Our studies indicate that for small networks (e.g., connecting up to

Y. Ben-Shimol (✉) · Y. Ben-Yehezkel · A. Dvir · M. Segal
Ben-Gurion University of the Negev, Beer Sheva, Israel
e-mail: benshimo@bgumail.bgu.ac.il

B. Ben-Moshe
College of Judea and Samaria, Ariel, Israel
e-mail: benmoshe@cs.bgu.ac.il

Y. Ben-Yehezkel
e-mail: yoav@yitran.com

A. Dvir
e-mail: azdvir@bgu.ac.il

M. Segal
e-mail: segal@cse.bgu.ac.il

ten base stations), the results obtained by human experts are adequate although they rarely exceed the quality of automated alternatives. However, the process of obtaining these results in comparison to automated heuristics is longer. In addition, when more base station sites need to be interconnected, the human approach is easily outperformed by our heuristics, both in terms of better results (fewer antennas) and in significant shorter calculation times.

Keywords Fixed-access wireless networks · Automated antenna positioning · Rural areas · Terrain preprocessing

1 Introduction

In wireless fixed access (WFA) networks a set of base stations (BSs) provide communication services to remote fixed users. Each base station antenna is planned to serve only its local potential subscribers and there is no need for continuous coverage of a given area by the set of BSs, particularly in rural areas. However, interconnections between the BSs are required in order to form a connected communication network where subscribers are able to communicate with each other. Here we assume that the interconnections between the base stations are obtained by positioning additional relay stations (RSs) in specific sites and using radio communication links. This assumption is valid for many rural areas where the sparsity of the population and the cost of wiring technologies motivate the use of wireless point-to-point links, and satellite based solutions provide insufficient bandwidth.

Installation and maintenance costs of the RSs motivate the minimization of the number of the RS sites. Despite the significant cost saving by eliminating even a single antenna site, an inefficient manual approach is used due to the computational complexity of the problem (Max-SNP-hard) (Feige 1996). Practical solutions to problems of antenna positioning of WFA networks are usually semi-automated. A professional designer, experienced in the essentials of installing WFA networks, suggests an antenna-positioning strategy. Interconnecting microwave antennas (as performed in this work) requires both Line Of Sight (LOS) and a distance limit between any two antennas. Connecting different antenna types may even succeed when LOS is broken, although path loss increases significantly (Blaunstein 1999), and is considered inappropriate for point to point communication. To the best of our knowledge, for a given rural area of 50 Km^2, a preliminary design has been obtained within a number of days, and by mixing human and computerized decisions. The final design is obtained after working over several weeks with much higher human involvement due to the lack of highly automated decision supporting tools. This information was given to us by experts from the industry planning such networks worldwide (see http://www.schema.com/; http://www.hexagonltd.com/; http://www.alvarion.com/).

It is natural that human design strategies carried out by experts usually achieve good results for small WFA networks. However as the number of BSs increases, the human approach becomes less efficient, more confusing and more time consuming.

An *informal* definition of the discussed problem follows: Given a terrain T and a set S of BSs, connect all BSs in that set by positioning the *smallest possible set* of RSs while considering terrain properties. An RS can connect any number of BSs/RSs

and a BS can also serve as an RS. This problem is referred to as CMBS (Connecting Multiple fixed access Base Stations).

Some solutions for automation of antenna-positioning problems in cellular networks have previously been suggested (Vasquez and Hao 2001; Anderson and McGeehan 1994; Tutschku 1998; Sherali et al. 1996; Hao et al. 1998). In addition, the CMBS problem is often raised in the context of minimizing the number of BSs that can serve all (or most) network clients (Ben-Moshe 2004). Solutions to the problem of serving a maximum surface of a geographical area with a minimum number of Base Transceiver Stations were also given (Calgeri et al. 2001).

The Steiner Minimal Tree Problem in networks (SMTP), where it is necessary to find a minimal cost tree spanning a given set of terminals with the possibility of using potential non-terminal nodes, resembles the CMBS problem. An unweighted graph, observing the triangle inequality for combinatorial distances, with the set of BSs serving as terminals, and the RSs as the non terminal nodes is a representation of the CMBS in terms of the Steiner problem. SMTP was proven to be an NP-hard intractable problem even in the special cases of Euclidean or rectilinear metrics. The Steiner problem has been thoroughly studied. An extensive summary of many related exact and approximated solutions and references are given in (Hwang et al. 1992; Prömel and Steger 2002). The best Steiner approximation ratio with polynomial time complexity was gradually improved from 2 to 1.55 in a series of papers (Takahashi and Matsuyama 1980; Zelikovsky 1993; Berman and Ramaiyer 1994; Zelikovsky 1996; Prömel and Steger 1997; Karpinsky and Zelikovsky 1997; Hougardy and Prömel 1999; Robins and Zelikovsky 2005). The best approximation algorithms discussed in (Hougardy and Prömel 1999; Robins and Zelikovsky 2005) are achieved by using iterative methods. However, as will be discussed later, the CMBS and the SMTP problems differ from the optimization perspective (e.g., target functions and the meaning of the approximation ratio of heuristics). Moreover, since the present work discusses a true real world problem with tens or hundreds of thousands of vertices, constants that are "hiding" behind the big $O(\cdot)$ notation for run time complexity become relevant when the algorithms are implemented.

For instances of the CMBS problem, Steiner approximation algorithms would try to minimize the following ratio

$$\frac{|B| + |S| - 1}{|B| + |S^*| - 1} > 1, \tag{1}$$

where B is the set of base stations, S^* is the (minimal) set of RSs participating in the optimal solution and S is the set of RSs participating in the approximate solution. $|\cdot|$ denotes the cardinality of the set (i.e., the number of nodes in the set). However, in this paper we are primarily concerned with reducing the number of RSs rather than minimizing the overall number of nodes. For a relatively large number of BSs, a small change in the number of required RSs may seem insignificant when using Expression (1). For example, if there are 50 BSs (which is a reasonable assumption) and two approximated solutions result with one and two RSs, respectively, the change may be considered insignificant when using Expression (1), but from the engineering point of view, the maintenance costs are twice as much. We try to minimize the

number of RSs and use the following expression as a measure of solution quality

$$\frac{|S|}{|S^*|} > 1. \tag{2}$$

The resemblance of the CMBS to SMTP motivated us to check ideas already implemented for the SMTP. We are particularly interested in the lowest running time complexity solution that can serve as building blocks for new fast iterative algorithms. Simple variations of minimal spanning trees and shortest paths algorithm are good candidates, due to their ease of implementation and short running times. Moreover, these methods have also been studied extensively in the context of facility location problems (Mitchell 2000; Agrawal et al. 1997; Har-Peled 1999; Tan 2004). In this paper we also suggest two novel heuristics based on a new grading scheme defined specifically for solving the CMBS problem. This will be discussed in more detail in Sect. 3.

To solve the design issues raised above, the goal of the present research is to establish a fully-automated antenna positioning algorithm for WFA networks, outperforming human approaches in terms of both selecting fewer RS antennas, and achieving this strategy in a shorter time. One should note that solutions of equal number of BSs serving the same number of subscribers could still be subjected to a significant difference from their CMBS point of view, that is, the number of RSs. Therefore, we believe that CMBS algorithms should be used as complementary objectives of applications that mix geographical information systems and radio propagation in rural environments (http://www.schema.com/; http://www.hexagonltd.com/). Usually these applications focus on optimization algorithms that attempt to minimize the number of antennas while providing service to the subscribers. However, interconnecting base station antenna sites directly or through a minimum set of relay stations is not considered in such applications.

From a theoretical aspect we show that the CMBS problem is *set-cover* hard to approximate. Section 2 presents a solution for connecting two BS antenna sites using the smallest set of RSs. Section 3 proves the NP-hardness of the CMBS problem and is also dedicated to heuristics that connect multiple (more than two) BSs, as well as to some examples of the CMBS solver application which we have developed. Section 4 provides details of the large scale experiment we conducted. Recommendations regarding future work are summarized in Sect. 5.

2 The point-to-point (P2P) case

This section discusses the simplest case of the CMBS problem: find a path with a minimal number of RSs connecting two remote BSs, one denoted as BS_s and the other as BS_t, where s, and t stand for source- and target-base station respectively. In spite of the "simplicity" of the problem, many fundamental issues arise and necessary intermediate steps need to be taken in this problem as well as in the more general case where more than two BSs need to be interconnected.

The problem of connecting two BSs with RSs is defined as follows: *Given a terrain T and exactly two base stations, connect these BSs using the smallest subset of potential relay stations. A BS can also serve as an RS.*

2.1 Terrain preprocessing

Since the proportion of the terrain models in our problem is great (a digital map covering an area of 100×100 Km^2 is not considered unusual), it is necessary to be preprocessed in order to reduce the input size by an order of magnitude before applying any facility location algorithm to it. To analyze whether two antennas that are positioned on two arbitrary points on the map can communicate directly, it is necessary to calculate the accumulated path-loss between these two points and compare it with the sensitivity threshold of the receiver (Blaunstein 1999). Two nodes directly connected by a radio link require a direct LOS and need to be in the proximity of 5–25 Km, depending on the antenna type and height, in order to guarantee that the cumulative path loss is small enough to allow a fast and reliable connection between the end points.

Calculating all pairs of points on a large-scale map is not feasible, therefore, we applied a preprocessing algorithm attempting to preserve the LOS property between potential antenna locations (usually positioned on top of hills or on high dominating places). We used a Digital Elevation Map (DEM) format, representing the map of the analyzed area as an elevation function of x, y coordinates, that is, $z(x, y)$. The amount of data given in a DEM map depends on its resolution and the size of the area under investigation. For a given area of 100×100 Km^2 with relatively coarse resolution of 100 m, there are approximately 10^6 potential sites resulting in $5 \cdot 10^{11}$ pairs. This is a large number for even a simple algorithm. Better resolutions (e.g., 10 m, 50 m) required for high quality solutions render the problem even more difficult to manage.

Two methods for finding positions of potential RSs were used. In the first a network of $N \times N$ grids was overlaid on the map, where N is a user-defined parameter. From each $N \times N$ grid, the two highest grid points among all the $N \times N$ points in that square are selected as potential antenna positions. The second method is based on a variation of using the map as an input for a convex hull algorithm. When a map serves as an input to a convex hull algorithm, the outcome consists of non-relevant facets below the surface. These non relevant facets are eliminated by padding the map margins with elevations that are lower than all the other points in the map; ensuring only a single facet below the surface of the convex hull. It is easier to eliminate this "artificial" low facet than to correct the output of the convex hull algorithm performed on the unpadded map. The positions of the potential RSs are then obtained from the list of coordinates of the convex hull nodes. Figure 1 presents the result of such preprocessing after removing the "artificial" facet and the additional points that had been added to the map.

There is no need for additional preprocessing when two other sites are defined for the base stations. This allows a single computation of a convex hull for each map and for later introducing the convex hull points as integral "implanted" candidate points in the input for the calculation of the CMBS problem. We wish to point out that although the method we suggested for finding potential RS positions using the $N \times N$ grids seems somewhat trivial, its strength is in its simplicity and its short running-time (which is a dominant factor in this research).

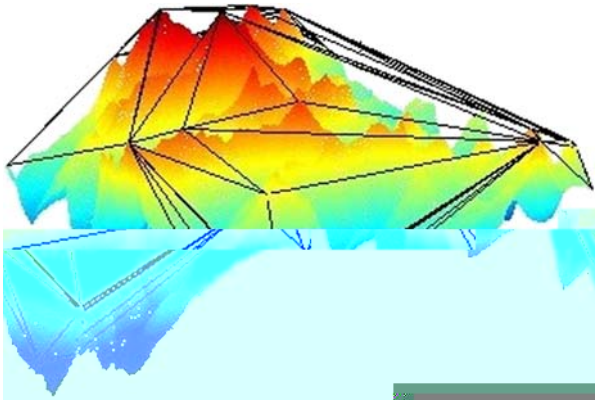

Fig. 1 Convex hull variation output on a map (after removing the single facet below the surface)

2.2 Positioning remote stations

Preprocessing the terrain with one of the methods described above allows us to construct a visibility graph. The visibility graph, denoted by $\mathrm{VG} = (V, E)$, where V is the set of all potential points for relay stations together with BS_s and BS_t. An edge $e = (v_i, v_j)$ denotes that the Euclidean distance between nodes v_i and v_j is smaller than a given threshold and there is a line of sight (LOS) between the nodes v_i and v_j. Both conditions are required to enable direct communication between v_i and v_j.

In the visibility graph VG all the shortest paths (in terms of number of edges) are given as suitable alternative solutions to select from, since, when testing these paths in the field, additional constraints such as restricted military zones, deep water reservoirs, swamps, etc., which cannot be used or are impractical for antenna positioning, are posed. Automation of selecting a feasible path from the set of the shortest paths that are returned by this process is not within the scope of the present research (although implemented for other researches with different scientific scopes).

The set of all shortest paths from BS_s to BS_t is discovered by a combination of the BFS and DFS search algorithms, which we refer to as Breadth First Depth Next (BFDN) algorithm.

The algorithm steps are as follows: Commencing at vertex BS_s, we transverse on the visibility graph VG using BFS and preserve edges of paths not previously visited. Unlike the ordinary BFS algorithm, upon reaching a node already reached with the same amount of hops, the edges of that path are also preserved. Upon reaching node BS_t by all nodes that could reach it with the same number of edges, the algorithm terminates. This yields a new graph G_{int} with many unnecessary edges of paths not reaching the destination and passing through nodes that may not appear in any of the shortest paths found leading to node t. In order to remove these edges, all the edges of G_{int} are reversed and another transversal begins from node BS_t using DFS and maintaining only the paths reaching node BS_s. This method produces a new graph G_{res} that contains all shortest paths from node BS_t to node BS_s. Reversing the edges of G_{res} yields the required solution. An illustration of the BFDN algorithm is given in Fig. 2.

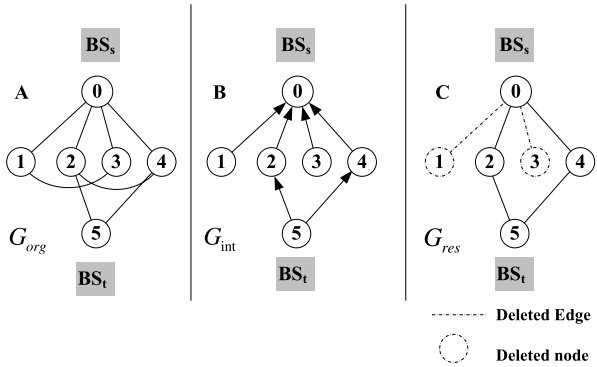

Fig. 2 Example to the BFDN algorithm

3 Establishing communication links among multiple base stations

In this section we deal with solutions to the more general problem, namely, connecting multiple BSs with a minimal set of RSs (i.e., CMBS). A formal definition of the CMBS problem in terms of graph theory is given below

Definition 3.1 (CMBS) Given an undirected graph $G = (V, E)$, a set $B = \{b_1, \ldots, b_n\}$ of n BSs and a set $R = \{r_1, \ldots, r_m\}$ of RSs such that $V = B \cup R$, find the smallest subset $R_{\text{opt}} \subseteq R$, such that the graph induced on $B \cup R_{\text{opt}}$ constructs a single connected component.

The intractability of the CMBS problem is proved to be NP-complete with a formal reduction from the general set cover problem to the CMBS problem. This reduction shows that the CMBS problem is "set cover hard". Several heuristics are developed and implemented in an attempt to resolve the CMBS problem and present solutions to one scenario using the application we developed for the current research.

The notation used in Definition 3.1 will serve us during the remainder of the paper. When discussing the solutions developed in this paper we will refer to the graph as VG (rather than G) in order to emphasize that the graph under discussion was obtained from the terrain pre-processing. That is, edges in VG connects nodes on the terrain with a direct LOS and in proximity of 5–25 Km, depending on the antenna type and height.

3.1 NP-completeness (MAX SNP-hard) proof

The CMBS problem, as shown below, is an optimization problem which has no constant approximation factor. That is, assuming that $P \neq NP$, for any suggested algorithm (A) solving the CMBS problem and a constant number C—there is always an input for which the ratio between the size of the solution (given by the algorithm A) and the size of the optimal solution is bigger than C. We argue that any instance of the set cover problem can be reduced to an instance of the CMBS problem such that an optimal solution for the CMBS problems implies an optimal solution for the set cover problem. Furthermore, every approximation for CMBS problem implies the same approximation for the set cover problem.

3.2 Heuristics for the CMBS problem

Since the CMBS is proved to be NP-hard, there is no tractable optimal solution, and

– **Step 4**: Determine a connected component S_k in S that is closest to T. Add that distance to cost. $T = T \cup \{S_k\}$, $S = S \backslash S_k$. $k = k + 1$. Go to **Step 3**.

The *T-MST* is a naive and fast algorithm. Note that the result stored in *cost* serves as an upper bound for the *number* (not positions) of required RSs, since the summation of distances may include RSs positioned on several shortest paths. Later this approximation is frequently used in the positioning heuristics given below as a grading mechanism enabling the comparison of the effectiveness of several possible solutions. We also use the fact that *T-MST* returns an upper bound Sect. 4, where we perform a comparison between heuristics using *T-MST* as one of the normalizing factors.

Regarding the lower bound of the *T-MST* algorithm, we observe that this algorithm achieves an optimal solution in the case when no reuse of RSs is required in order to connect all BSs. In other words, the paths connecting different BSs are all disjoint. This is due to the fact that our algorithm works in a Kruskal-like fashion. On the other hand, a lower bound of a diameter (in terms of number of RSs) defined by the two farthest BSs in the visibility graph VG is applicable in the general case. We notice that remaining algorithms presented in this section have better performance than the *T-MST* algorithm.

In what follows we present a set of heuristics for the CMBS problem which are concerned with the positioning of RSs.

Simple Minimal Spanning Tree (*S-MST*): the *S-MST* heuristic is the first heuristic that allows us to actually position RS antennas on the given terrain and that has three variations:

1. First, select a predefined initial BS. Then calculate the shortest path to its nearest neighbor. If several such paths are available, pick the first one on the list. In each step henceforth, connect the first BS closest to the already connected component (any already connected RS or BS).
2. The same as variant 1 with the following distinctions: instead of predefining the initial BS, select it randomly, and if more than one shortest path to the closest connected component exists, select one of these shortest paths randomly.
3. This variation resembles the second variation; however, rather than connecting the entire shortest path to the already connected component, add only a single RS to the solution until all BSs are connected.

The first two variants of the *S-MST* heuristic are similar to the SPH heuristic used to solve the SMTP (see Hwang et al. 1992; Prömel and Steger 2002). The present variant differs mostly in the way the shortest path is found; we use the BFDN algorithm discussed in Sect. 2.2. This is important for more advanced algorithms and will be discussed later. The third variant attempts to avoid the preference given to shortest paths by adding only the first RS on the shortest path. The *S-MST* heuristic provides fast solutions and is used later by some more advanced heuristics given below.

Random Batch (*RB-MST*): Simply apply the *S-MST* heuristic several times, each time the initial BS, one of the shortest paths or the added RS being selected at random. *RB-MST* returns the best results.

We now wish to describe *T-MST*, *S-MST* and *RB-MST* on the VG example given in Fig. 4. We begin by examining the *T-MST* approximation. The shortest path distance between BS_a and BS_b is three, and the shortest path distance from BS_c to BS_a

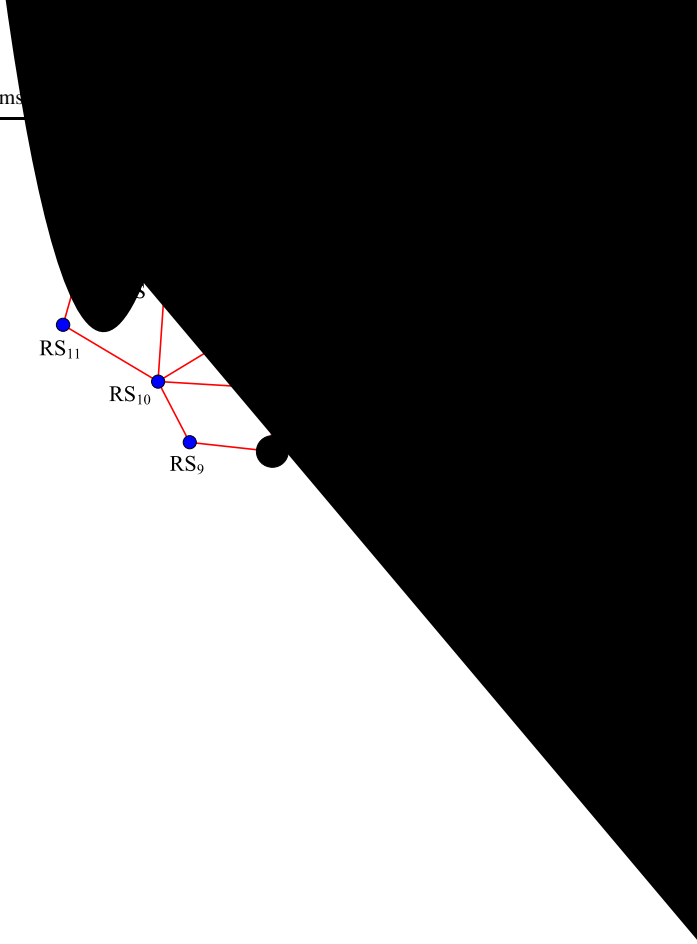

Fig. 4 Visibility graph VG with three BSs (*big circles*) and 11 RSs (*small circles*) as for demonstrating the operation of the positioning algorithms

is usually much smaller, an immediate motivation is to use only these independent shortest paths. This leads to a faster implementation of *GI-MST* where the grading function is carried out only on the RS-min-cut set between any two BSs on VG.

Greedy and Random (*GR-MST*): run *GI-MST* several times, selecting the *GI-MST* parameters at random. The initial BS, which RS-min-cut is set to select (if there are several of the same size), and the destination BS (if several shortest paths reach more than one BS), are examples to the random selection that can be done. In general, whenever a selection is made, it is done randomly. *GR-MST* returns the best results.

The implications of *GI-MST* and *GR-MST* are demonstrated in Fig. 4. Assuming that we begin at BS_a, there are five shortest paths to BS_b. The size of the RS-min-cut set is two, resulting from any set of nodes whose distance to the BS is the same. We select the first RS-min-cut set (RS_5 and RS_9). Then RS_5 is treated as a BS, and *T-MST* is recalculated. Following this order *T-MST* returns four, whereas starting RS_9 *T-MST* returns five. Note that if *S-MST* is used in place of *T-MST*, the result would have been optimal. Figure 4 also enables examining the set of random selections the *GR-MST* makes. For example, the algorithm may start at any of the three BSs. There are five shortest paths from BS_a to BS_b, but also to BS_c. There are three RS-min-cut sets of size two on any selected shortest path. We do not obtain better results, but for larger and non-symmetrical inputs, *GR-MST* occasionally produces better results than *GI-MST*.

3.4 Relay stations grading (RSG) heuristics

The algorithms discussed in the previous section are all based on RSs that are on shortest paths between BSs. However, better solutions could be obtained by using RSs that are not on any shortest path between BSs. Here we present another class of heuristics based on the concept of grading RSs, where a grade of an RS should indicate how likely it would be included in a good solution of the CMBS problem.

The following two heuristics are based on the fact that the optimal solution might contain an RS which does not reside on any of the shortest paths from any BS to any other (closest) BS (see Fig. 5). This infers that none of the earlier suggested MST-based heuristics will produce an optimal solution. A naive idea is to run a *GI-MST*-like algorithm on all the RSs sets as potential candidates. Although it solves the example in Fig. 5, it is too slow and impractical for real-life large input scenarios, and in most cases achieves poor results.

Basic Relay Station Grading (*B-RSG*): *B-RSG* enables locating RSs that are not necessarily on any shortest path (between two BSs). *B-RSG* grades candidate RSs

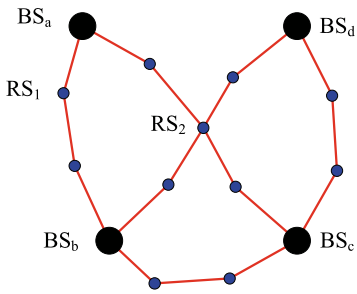

Fig. 5 An example with four BSs (*big circles*) and 11 RSs (*small circles*) where the optimal solution is five (using the *RS at the center of the graph*), while any solution based on the shortest path-like methods is six (using only *outer RSs* that are on shortest paths between BSs)

based both on their reuse property and on their path distances. The fact that the number of paths on which the RS can be used improves its grade and allows it to break loose from shortest path based solutions. Several grading functions were tested and compared. The result of this comparison is a grading scheme that provides the required effect.

For each connected component of VG[B] (that is, each base station BS_i, or a set of BSs, directly connected without RSs), find the closest BS', denoted by $d_1 = d(\mathrm{BS}_i, \mathrm{BS}')$, where d is the link distance in terms of RS hops. For each RS, such that $d(\mathrm{RS}, \mathrm{BS}_i) \leq d(\mathrm{BS}_i, \mathrm{BS}')$, we find the closest BS'' (excluding BS_i). Denote $d_2 = d(\mathrm{RS}, \mathrm{BS}_i) + 1$ and $d_3 = d(\mathrm{RS}, \mathrm{BS}'')$. The grade for the RS using BS_i is given by the expression: $g(\mathrm{RS}, \mathrm{BS}_i) = f(d_1, d_2, d_3)$, where $f(d_1, d_2, d_3) = (d_1/(d_2 + d_3))^c$ with $1 \leq c \leq 2$. The total grade of an RS is $g(\mathrm{RS}) = \ {}_{i=1}^{n} g(\mathrm{RS}, \mathrm{BS}_i)$ where n is the number of connected components of BSs in VG[B].

.96 0 T03 Tc 626 0 0 9.9626 293.27413 471.984 Tm [(RS)] TJ ET 0 Tc BT /F7 1 Tf 9.9626 0 0

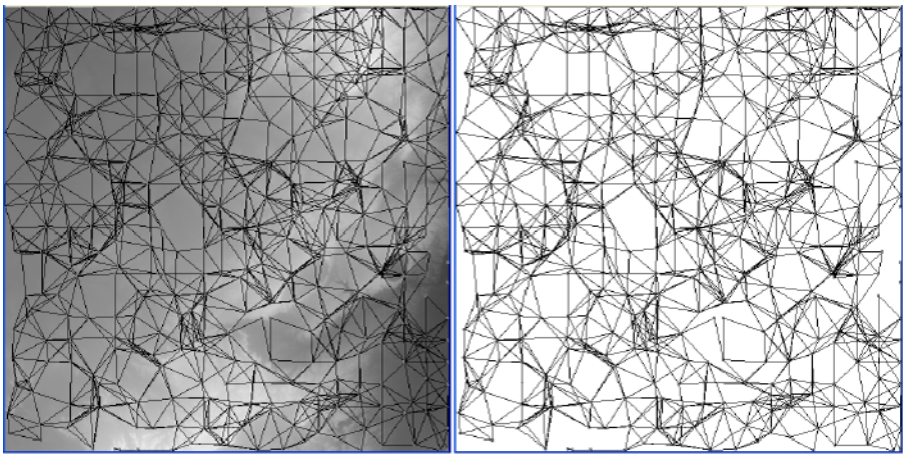

Fig. 6 VG of 450 RS (with 10 m antennas heights and 10 Km transmission range). On the *left*: VG displayed over 100 × 100 Km^2 DEM (*lighter grey* refer to higher altitudes). On the *right*: the same VG displayed with the DEM hidden

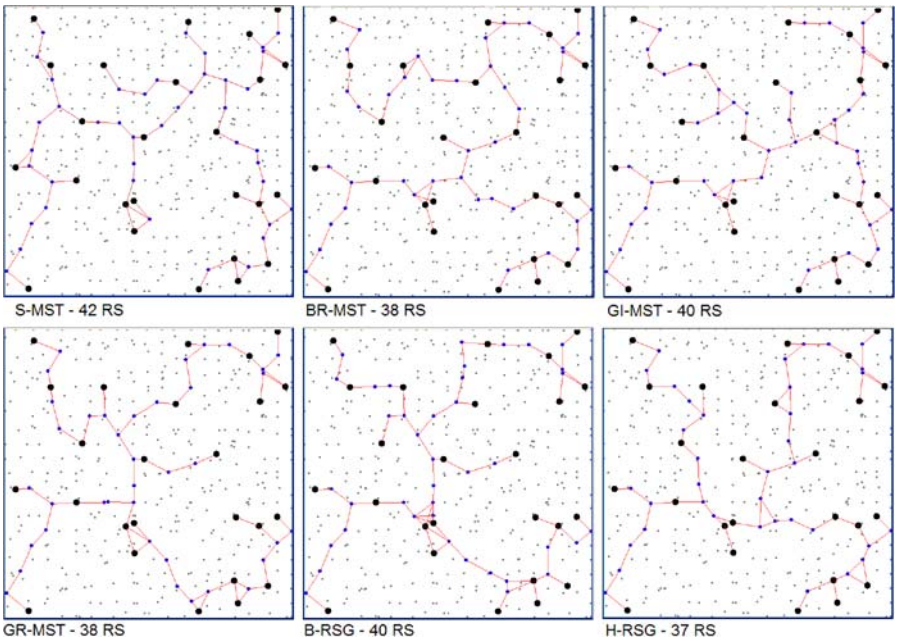

Fig. 7 An example of a single scenario solved by all the suggested heuristics (the T-MST value of this scenario is 46)

with all antennas height 30 m, and transmission ranges 10 km, with 1,000 potential RSs).

Many different possible sets of scenarios were tested. The experiment consisted of a total of approximately 7,000 scenarios (about one quarter of the possible permu-

Table 1 List of tested parameters in the large-scale experiment

Parameter	Permutations
Maps	17
Total number of BSs on a map (10/25/50/100 BSs)	4
Number of BSs positions for each parameter vector	10
Antenna Heights (10/20/30/50 m)	4
Transmission ranges (5/10/20 km)	3
Potential relay stations (500, 1000, 2500, 10000 relays)	4

tations) using all the heuristics detailed in Sect. 3. Non-realistic cases, such as long ranged antennas used with many RSs, or short ranged antennas with few potential RSs, were eliminated.

17 high-resolution elevation maps (100×100 km^2), each representing a different type of terrain such as flat, hills, dunes, mountains, lakes, etc., were used. We positioned the selected number of BSs on the map either by selecting a subset of the potential RS (which is a reasonable strategy, as BS are usually positioned on high altitude points) or at random. For each permutation of the parameters (i.e., number of BSs, number of potential RSs, antenna heights and transmission ranges) we tested ten different positions for BSs.

4.2 Results

In order to compare the results obtained by all the heuristics a useful comparison method was implemented, titled the *Best Of All* (*BOA*). For each scenario *BOA* picks the best result obtained by any of the heuristics. Using this method we can compare results obtained by each heuristic to the minimal number of RSs obtained by any of the other heuristics (i.e., comparing it to the best performance obtained for the specific scenario).

Figure 8 presents a bar-chart showing the mean number of RSs required to connect 10, 25, 50 and 100 BSs. The graph compares between results obtained by each heuristic and the results are presented with an increasing order of the mean effectiveness of the algorithms (i.e., a decreasing order of number of RSs). The results indicate that *H-RSG* was the most effective heuristic and is only slightly less effective than the *BOA*. *B-RSG* almost always produced the best results compared to all the "single iteration" heuristics (*T-MST*, *S-MST* and *GI-MST*). However, *B-RSG* also has the longest run-time among the other "single iteration" based heuristics. It takes the same time to run *B-RSG* once or to run *RB-MST* many times and usually obtains better results, due to the randomization embedded in *RB-MST*.

A reasonable assumption is the decrease in the average number of required RSs per BS while increasing the number of BSs. The rational behind this assumption is the decrease in the average distance between base stations and also because BSs also serve as RSs. This assumption was found accurate, using all heuristics, as can be seen in Fig. 9.

Variation of heuristics: We examined which factors cause maximal variation between heuristics. The two dominating factors were found to be the number of BSs

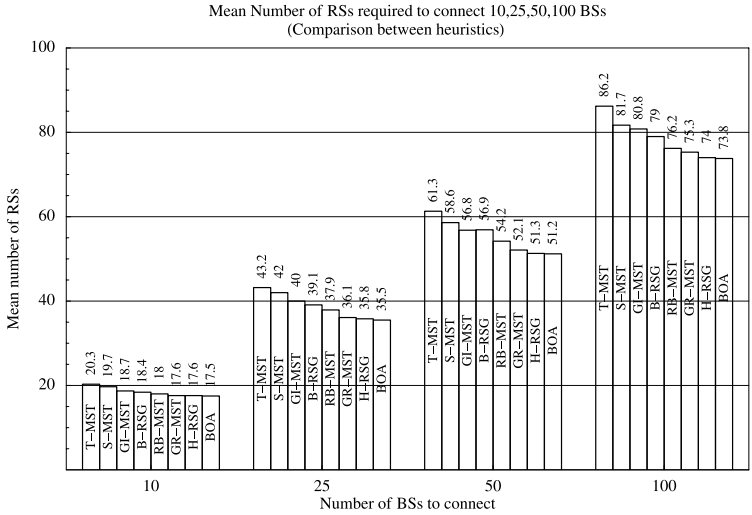

Fig. 8 Mean number of RSs required to connect 10, 25, 50 and 100 BS—comparison between heuristics

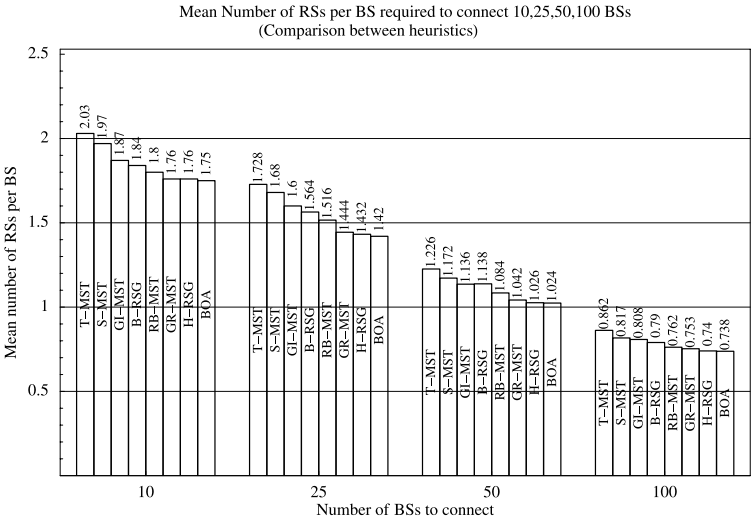

Fig. 9 Mean number of RSs required per BS when connecting 10, 25, 50 and 100 BSs

needed to be connected and the number of edges on VG. The latter parameter is related on the one hand to the number of possible solutions, thus increasing the variation between heuristics, and on the other hand reduces the number of RSs required for a solution, thus decreasing the variation between heuristics. Figure 10 shows that a smaller variation is expected when VG is denser, and a larger variation is expected when more BSs are required to be interconnected.

Comparison among heuristics: To compare among heuristics we calculated the normalized effectiveness of each heuristic solution for each scenario. Normaliza-

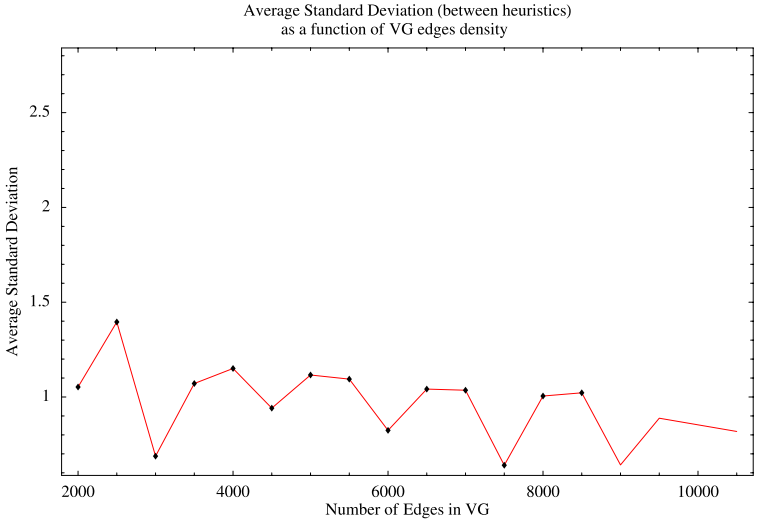
Average Standard Deviation (between heuristics)
as a function of VG edges density
Average Standard Deviation
2.5
2
1.5
1
2000
4000
6000
8000
10000
Number of Edges in VG

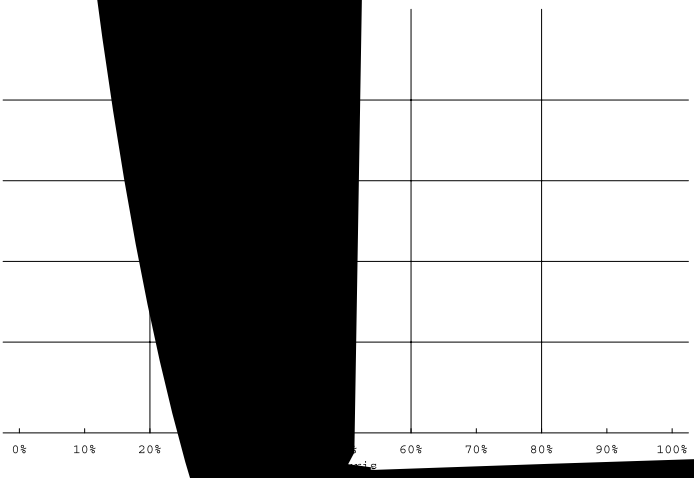
0%
10%
20%
60%
70%
80%
90%
100%

Table 2 Average run time of the given algorithms for different scenarios

Algorithm	Number of BSs/RSs			
	10/500	25/1000	50/2500	100/5000
S-MST	1	8	52	140
GI-MST	35	107	287	532
B-RSG	72	183	530	1236
RB-MST	402	1255	2743	5943
GR-MST	393	1214	2717	5901
H-RSG	424	1250	2813	6012

5 Future research

Further research may include additional heuristics and fine-tuning techniques to the ones suggested here. A generalization of the CMBS problem to a weighted CMBS problem in which the objective function is weighted according to certain parameters of the relay stations (height, maintenance costs, etc.), subject to the same constraint of connecting all BSs seems a good direction for research. One obvious application for this generalization is to solve the CMBS problem while minimizing the installation cost, for example by assuming different costs to different RS types. We also believe that the algorithm for selecting potential RS antenna positions can be improved in terms of runtime complexity and approximation factor. Another aspect of the CMBS problem is to compute a lower bound for a given input by employing continuous linear programming based methods.

References

Agrawal, P.K., Har-Peled, S., Sharir, M., Varadarajan, K.R.: Approximate shortest paths on a convex polytope in three dimensions. J. ACM **44**, 567–584 (1997)

Anderson, H.R., McGeehan, J.P.: Optimizing microcell base station locations using simulated annealing techniques. In: Proc. of the IEEE Vehicular Technology Conference, pp. 858–862 (1994)

Available at http://www.alvarion.com/

Available at http://www.hexagonltd.com/

Available at http://www.schema.com/

Ben-Moshe, B.: Geometric facility location optimization, Ph.D. Thesis, Ben-Gurion University of the Negev (2004)

Berman, P., Ramaiyer, V.: Improved approximations for the Steiner tree problem. J. Algorithms **17**, 381–408 (1994)

Blaunstein, N.: Radio Propagation in Cellular Networks. Artech House, Boston (1999)

Calgeri, P., Kuonen, P., Nielsen, F.: Combinatorial optimization algorithms for radio network planning. Theor. Comput. Sci. **263**, 235–245 (2001)

Efroymson, M.A., Ray, T.L.: A branch and bound algorithm for plant location. Oper. Res. **14**, 361–368 (1966)

Feige, U.: A threshold of $\ln n$ for approximating set cover (preliminary version). In: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing, pp. 314–318 (1996)

Hao, J.K., Dorne, R., Galinier, P.: Tabu search for frequency assignment in mobile radio networks. J. Heuristics **4**, 47–62 (1998)

Har-Peled, S.: Constructing approximate shortest path maps in three dimensions. SIAM J. Comput. **28**, 1187–1197 (1999)