Practical Approaches to Partially Guarding a Polyhedral Terrain

Frank Kammer[1], Maarten Löffler[2], Paul Mutser[2], and Frank Staals[2]

[1] Institut für Informatik, Universität Augsburg, 86135 Augsburg, Germany
kammer@informatik.uni-augsburg.de
[2] Department of Information and Computing Sciences, Utrecht University
{m.loffler, f.staals}@uu.nl, paulmutser@gmail.com

Abstract. We study the problem of placing guard towers on a terrain such that the terrain can be seen from at least one tower. This problem is important in many applications, and has an extensive history in the literature (known as, e.g. multiple observer siting). In this paper, we consider the problem on polyhedral terrains, and we allow the guards to see only a fixed fraction of the terrain, rather than everything. We experimentally evaluate how the number of required guards relates to the fraction of the terrain that can be covered. In addition, we introduce the concept of dominated guards, which can be used to preprocess the potential guard locations and speed up the subsequent computations.

1 Introduction

Terrains are a key concept in Geographic Information Science. They are the topic of interest in may different problems, ranging from determining how water flows along a terrain [16,22] to computing valleys and ridges [19]. We study the problem of guarding a terrain; that is, we wish to place a small number of guards such that they can together can see the terrain. The applications for this problem are numerous. Consider for example protecting the border of a country, placing watchtowers to protect against forest fires [6], or determining where to place base stations for a telecommunication network [8,24]. See also Floriani and Magillo [11] for an extensive treatment of the subject.

There are two standard representations for terrain data. We can store a terrain as a digital elevation model (DEM), which is a two-dimensional grid with height values, or as a *polyhedral terrain*: a planar subdivision—usually a triangulation—in which each vertex has an associated height. Heights are linearly interpolated in the interior of a face. This yields a polyhedral surface in $\mathbb{R}^3$.

For grid-based terrains, guarding is well-studied. Franklin *et al.* [12] present a greedy approach, Kim *et al.* [18] investigate heuristics for placing guards, and Zhang and Lu [20] use improved simulated annealing to determine where to place the guards. However, grid terrains are less suited for visibility problems than polyhedral terrains [4,23]. Furthermore, polyhedral terrains allow a more compact representation of the data, which may allow us to avoid heuristics when working in external memory [21].

Informally speaking, the viewshed of a guard is the part of the terrain that it can see. An example is shown in Fig. 1. Computing viewsheds of guards at fixed locations is itself useful in many applications, e.g. bird behavioral studies [3,5]. There are efficient algorithms to compute the viewshed for a given guard [17], and even computing the part of the terrain visible by a set of guards can be done efficiently [14]. Unfortunately, it is NP-hard to determine *where* to place a minimum number of guards such that they can together see the entire terrain [7]. Moreover, Eidenbenz *et al.* [10] showed that there is no polynomial time algorithm that can approximate the number of guards required to cover the whole terrain consisting of n triangles within a factor $O(\ln(n))$ of the optimum unless some complexity classes collapse, which is very unlikely.

Partial Covers. The digital model of the terrain that we work with is often imprecise, and even if we have the true heights for all points on the terrain, there other factors, such as vegetation, that impact visibility. So, instead of requiring that the guards see the entire terrain it may be sufficient if they see a large portion of the terrain. This raises the question if we can efficiently solve or approximate this so-called ε-guarding problem.

To define the ε-guarding problem precisely, we need some definitions. A *guard* g is a point (tower) at height h above a polyhedral terrain $\mathcal{T}$. It can see, or *cover*, a point $p \in \mathcal{T}$, if the open ended line segment $\overline{gp}$ lies entirely above the terrain. We also say that p is *visible* from g. The maximal set of points on $\mathcal{T}$ that is visible from g, denoted $\mathcal{V}(g)$, is the *viewshed* of g. The viewshed of a set of guards $\mathcal{G}$ is the maximal set of points on $\mathcal{T}$ visible from at least one guard in $\mathcal{G}$, that is, $\mathcal{V}(\mathcal{G}) = \bigcup_{g \in \mathcal{G}} \mathcal{V}(g)$. We can measure the *size* $[\![\mathcal{T}']\!]$ of a part $\mathcal{T}'$ of $\mathcal{T}$, and thus the size of a viewshed, in two ways. Either we consider the *area* of $\mathcal{T}'$, or we consider the *number of terrain vertices* in $\mathcal{T}'$. Our algorithms can be used in both cases—only the running time changes through a different viewshed computation.

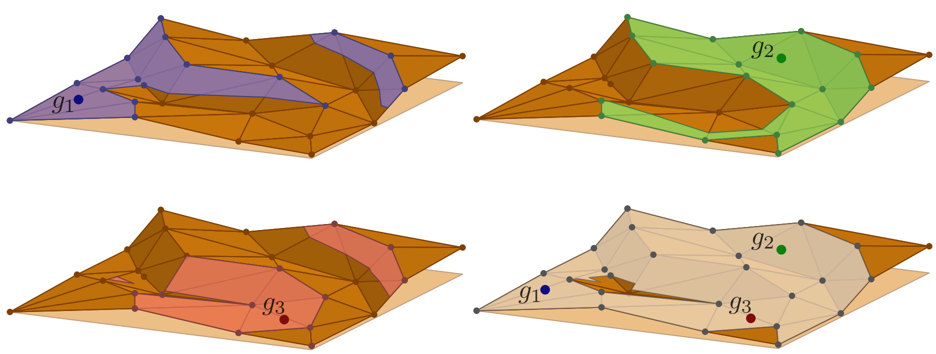

Fig. 1. The viewshed $\mathcal{V}(g_i)$ for each guard g_i , and the viewshed $\mathcal{V}(\{g_1, g_2, g_3\})$.

Definition 1 **(ε-cover, ε-guarding problem).** *Given a polyhedral terrain $\mathcal{T}$, a height h, and a value $\varepsilon > 0$, an ε-*cover *of $\mathcal{T}$ is a set of* guards *$\mathcal{G}$ that can together see at least a fraction of $(1-\varepsilon)$ of the terrain, i.e., a set of* guards *for which $[\![\mathcal{V}(\mathcal{G})]\!] \geq (1-\varepsilon)\,[\![\mathcal{T}]\!]$. In the ε-*guarding problem*, the goal is to find a minimum sized ε-cover.*

Our Contribution. In Section 2, we generalize Eidenbenz *et al.*'s [10] result and show that the same lower bound holds for ε-guarding problem for any $\varepsilon \leq 1 - 1/n^x$ with $x < 1$. Next, we make two main observations in this paper, which we experimentally validate. In Section 3 we analyse a practical approach to compute an ε-cover: a greedy algorithm that simply places the guard that covers the largest possible unguarded area, and continues to place more guards until a desired fraction of the terrain is covered. By using ideas from the approximation techniques known for the so-called Set-Cover problem [15] we show that the number of required guards can be related to the optimal number of guards required to cover the whole terrain. We also implemented the approach, and show that in real-world terrains, the number of required guards to cover 95% of a mountainous region of roughly 150km^2 is typically less than 20.

In Section 4, we introduce the concept of *dominating guards*: a potential guard tower *dominates* another potential guard tower if it can see at least the same part of the terrain. Since dominated guards are never necessary in an optimal solution, the computation of a good set of guards can be sped up by precomputing and deleting the dominated guards. We can also relax the concept and say a guard is dominated if another guard sees *almost* everything it sees. We show experimentally that in real-world terrains, computing dominated guards typically reduces the problem size by 15-40%, depending on the resolution of the terrain. For the relaxed version, this percentage drops drastically, reducing the problem size to as little as 10% of the original, but comes at the cost of potentially not allowing every solution anymore. However, we show that in practice, the greedy approach still finds a solution of the same size as without preprocessing.

The idea of reducing a problem's input size by transforming it to another instance, while preserving theoretical guarantees on the solution, is common in theoretical computer science. Agarwal *et al.* [1] introduce the concept of *core-sets*: a subset P' of a set of objects P (clasically, points in a d-dimensional space) with the property that for some function f that one is interested in, $f(P')$ differs from $f(P)$ by at most a factor $(1-\varepsilon)$. A related concept in complexity theory is *kernalization* [9].

2 Lower Bound

In this section we show a lower bound on the computational complexity of the ε-guarding problem. We first focus on the case where we measure viewsheds and terrain sizes $[\![\mathcal{T}']\!]$ with respect to the number of visible terrain vertices. Eidenbenz *et al.* [10] have shown in Lemmata 6 and 7, which are part of the proof of the non-approximability result for guarding-terrain problem, that an instance I for the so-called Set-Cover problem can be transformed into an instance I' for the guarding problem such I has a solution of size k if and only if I' has a solution of size $k+4$. Because of this and because optimal solutions of the Set-Cover problem can be of arbitrarily large size k we only need to show a lower bound for the so-called α-weak Set-Cover problem and obtain the same lower bound for the ε-guarding problem where $\varepsilon = 1-\alpha$.

Definition 2. *Given a number $0 \leq \alpha \leq 1$ and a tuple $(U, \mathcal{C})$, where U is a finite set called* universe *and $\mathcal{C}$ is a collection of subsets of U with $\bigcup_{X \in \mathcal{C}} X \subseteq U$, an α-weak* set cover *for $(U, \mathcal{C})$ is a collection $S \subseteq \mathcal{C}$ such that the union of all sets in S contains at least $\alpha|U|$ elements. In the special case $\alpha = 1$, we also call S a* set cover *for $(U, \mathcal{C})$. The* size *of S is its cardinality. In the (α-weak)* Set-Cover *problem, we have to find an (α-weak) set cover of minimal size.*

Lemma 1. *The α-weak* Set-Cover *problem has no polynomial-time approximation algorithm of ratio $c \ln n$ unless $\mathrm{P} = \mathrm{NP}$, where n is the size of the universe, $\alpha \geq 1/n^d$ with $d < 1$, and $c > 0$ is an appropriately chosen constant.*

Proof. Alon, Moshkovitz, and Safra [2] showed that, for an appropriately chosen constant $c'' > 0$, there is no polynomial-time approximation algorithm of ratio $c' \ln |U'|$ for the Set-Cover problem with universe U' unless $\mathrm{P} = \mathrm{NP}$. Since we can figure out in polynomial time if a set-cover instance has no solution, there can not exist a polynomial-time approximation of approximation ratio $c' \ln |U'|$ for the Set-Cover problem restricted to solvable instances.

We now show the lemma by a reduction from the Set-Cover problem to the α-weak Set-Cover problem. Let $(U', \mathcal{C}')$ be a solvable Set-Cover instance. Take $k = \lceil 1/\alpha \rceil$, $U = U' \cup \{x_1, \ldots, x_{k|U'|}\}$ where $x_1, \ldots, x_{k|U'|}$ are new elements, and $\mathcal{C} = \mathcal{C}' \cup \{\{x_1\}, \ldots, \{x_{d|U'|}\}\}$.

Note that each solution of the set-cover instance $(U', \mathcal{C}')$ is also a solution for the α-weak-set-cover instance $(U, \mathcal{C})$. For the reverse direction, assume that there is a solution $\mathcal{S} \subseteq \mathcal{C}$ for the α-weak-set-cover instance, and let $V := \bigcup_{X \in \mathcal{S}} X$ be the elements covered by $\mathcal{S}$. Intuitively speaking, for each element of U' that is not covered, there is an element of $U \setminus U'$ covered by the solution. More exactly, $|U' \setminus V| \leq |(U \setminus U') \cap V|$. To obtain a solution $\mathcal{S}'$ for the set-cover instance $(U', \mathcal{C}')$ from $\mathcal{S}$, for each $u \in U' \setminus V$, we replace a set in $\mathcal{S} \cap (\mathcal{C} \setminus \mathcal{C}')$ by a set $C' \in \mathcal{C}'$ with $u \in C'$. Note that set C' must exist since $(U', \mathcal{C}')$ has a solution.

To sum up, each solution of $(U', \mathcal{C}')$ can be turned into a solution of the same size for $(U, \mathcal{C})$ and vice versa. Then $(U, \mathcal{C})$ has no polynomial-time approximation algorithm of ratio $c' \ln |U'| = c' \ln(|U|/(k+1)) \leq c' \ln(n/(2n^d)) \leq c'(1-d)\ln(n/2) \leq c \ln n$ if c is chosen appropriately. □

Corollary 1. *The ε-guarding problem has no polynomial-time approximation algorithm of ratio $c \ln n$ where n is the size of the universe, i.e., the number of the terrain vertices, $c > 0$ is an appropriately chosen constant, and $\varepsilon = 1 - \alpha \leq 1 - 1/n^d$ with $d < 1$.*

If we want to measure terrain sizes $\llbracket \mathcal{T}' \rrbracket$ with respect to the visible area, we can first decompose the terrain into maximal regions such that each region is visible from the same set of guards. We then take the size of the universe to be the number of such regions. If we have $m \leq n$ (potential) guards, there are at least $\Omega(m^2n^2)$ and at most $O(m^3n^2)$ such regions, and they can easily be computed in polynomial-time [14]. Thus, also in the case with terrain sizes measured by the visible area, we obtain a $\Omega(\ln n)$ lower bound on the approximation ratio of an ε-cover, for small enough ε.

Note that, if $\varepsilon = 1-1/n^x$ with $x = 1$, we need to see one vertex of the terrain to solve the vertex viewsheds ε-guarding problem. Clearly, such a solution can be found easily.

3 Greedy Approach

3.1 Algorithm and Analysis

Consider the following simple and straightforward algorithm GreedyGuard that given a terrain $\mathcal{T}$, a parameter ε, and a set of potential guard locations $\mathcal{P}$, computes an ε-cover of $\mathcal{V}(\mathcal{P})$. So, when the guards in $\mathcal{P}$ can together see the entire terrain $\mathcal{T}$, GreedyGuard computes an ε-cover of $\mathcal{T}$.

Algorithm GreedyGuard$(\mathcal{T}, \varepsilon, \mathcal{P})$

1. Compute the viewsheds for all guards in $\mathcal{P}$.
2. Let $\mathcal{G} = \emptyset$ and $\mathcal{R} = \mathcal{P}$.
3. **while** $\llbracket \mathcal{V}(\mathcal{G}) \rrbracket \leq (1-\varepsilon) \llbracket \mathcal{V}(\mathcal{P}) \rrbracket$ **and** $\mathcal{R} \neq \emptyset$ **do**
4. Select a guard $g \in \mathcal{R}$ that maximizes the size $\llbracket \mathcal{V}(g) \setminus \mathcal{V}(\mathcal{G}) \rrbracket$, i.e., the size of the region it can cover but is not covered by $\mathcal{G}$ yet.
5. Remove g from $\mathcal{R}$ and add it to $\mathcal{G}$.
6. **return** $\mathcal{G}$

We now show that the selected set of guards $\mathcal{G}$ has size at most $O(k/\varepsilon)$, where k is the number of guards required in an optimal solution to cover $\mathcal{V}(\mathcal{P})$.

Lemma 2. *Let $\mathcal{T}$ be a terrain, let $\mathcal{P}$ be a set of potential guard locations, and let $\varepsilon \in (0,1]$.* GreedyGuard *computes an ε-cover of $\mathcal{T}' = \mathcal{V}(\mathcal{P})$ of at most $O(k/\varepsilon)$ guards, where k is the size of an optimal 0-cover of $\mathcal{T}'$.*

Proof. Consider the guards $g_1, .., g_\ell$ picked, in that order, by the greedy algorithm, and let ρ_i denote the fraction of $\mathcal{T}'$ that remains uncovered by the first i guards, that is, $\rho_i = \llbracket \mathcal{R}_i \rrbracket / \llbracket \mathcal{T}' \rrbracket$, where $\mathcal{R}_i = \mathcal{T}' \setminus \bigcup_{j=1}^{i} \mathcal{V}(g_i)$.

The greedy algorithm stops once the fraction of $\mathcal{T}'$ that remains uncovered is at most ε. That is, $\rho_\ell \leq \varepsilon$. We claim that this is the case for $\ell = ck/\varepsilon$, for some $c \in \mathbb{R}$.

For any i, with $0 \leq i < \ell$, the remaining part $\mathcal{R}_i$ of $\mathcal{T}'$ can be covered with k guards. So, there is a guard, say $g^* \in \mathcal{P}$, that covers at least $\llbracket \mathcal{R}_i \rrbracket / k$. It follows the next guard g_{i+1} picked by the greedy algorithm covers at least that much. Thus, we have $\rho_{i+1} [\![\mathcal{T}']\!] = [\![\mathcal{R}_{i+1}]\!] \leq [\![\mathcal{R}_i]\!] - [\![\mathcal{R}_i]\!] / k = \rho_i [\![\mathcal{T}']\!] - \rho_i [\![\mathcal{T}']\!] / k$, and therefore $\rho_{i+1} \leq \rho_i - \rho_i/k = \rho_i(1 - \frac{1}{k})$. It follows that $\rho_i \leq (1 - \frac{1}{k})^i$.

We thus have to show that $\rho_\ell \leq (1 - \frac{1}{k})^\ell \leq \varepsilon$. In case $k = 1$ this is trivially true. In case $k \geq 2$, we have that $(1 - \frac{1}{k})^k \leq \frac{1}{e}$. So, we have to show that

$$\left(1 - \frac{1}{k}\right)^\ell = \left(1 - \frac{1}{k}\right)^{ck/\varepsilon} = \left(\left(1 - \frac{1}{k}\right)^k\right)^{c/\varepsilon} \leq (1/e)^{c/\varepsilon} = 1/(e^{c/\varepsilon}) \leq \varepsilon.$$

Using basic calculus we can reduce $1/(e^{c/\varepsilon}) \leq \varepsilon$ to $\varepsilon \ln(1/\varepsilon) \leq c$. This holds for all $c \geq \frac{1}{e}$. The lemma follows. □

We now analyze the running time of this algorithm. If we are in the case where terrain sizes are measured by the number of terrain vertices, it is easy to compute the viewshed of a guard g in $O(n^2)$ time by simply checking if the line segment between g and each terrain vertex lies above the terrain. Thus, if we are given m potential guards in $\mathcal{P}$, this takes $O(mn^2)$ time in total. From Lemma 2 it follows that the algorithm selects at most $O(k/\varepsilon)$ guards. Finding a guard $g \in \mathcal{R}$ that maximizes $[\![\mathcal{V}(g) \setminus \mathcal{V}(\mathcal{G})]\!]$ takes $O(mn)$ time, so selecting all $O(k/\varepsilon)$ guards takes $O(kmn/\varepsilon)$ time in total. Thus, our algorithm can be implemented in $O(mn^2 + kmn/\varepsilon)$ time.

We can improve on this by preprocessing our terrain for visibility queries. In $O(n\alpha(n) \log n)$ time, where α is the extremely slow growing inverse of Ackermann's function, we construct a data structure that can answer visibility queries from a fixed viewpoint in $O(\log n)$ time [7]. Using this data structure we can compute all viewsheds in $O(mn\alpha(n) \log n + mn \log n) = O(mn\alpha(n) \log n)$ time. Thus, we conclude:

Theorem 1. *Let $\mathcal{T}$ be a terrain, let $\mathcal{P}$ be a set of potential guard locations, and let $\varepsilon \in (0, 1]$.* GreedyGuard *computes an ε-cover of $\mathcal{T}' = \mathcal{V}(\mathcal{P})$ of at most $O(k/\varepsilon)$ guards in $O(mn\alpha(n) \log n + kmn/\varepsilon)$ time, where k is the size of an optimal 0-cover of $\mathcal{T}'$, and m is the number of potential guards in $\mathcal{P}$.*

We remark that the GreedyGuard algorithm is essentially the well-known greedy algorithm for the Set-Cover problem. Hence, our results immediately transfer to the Set-Cover problem as well. Indeed, when we compute the required guards/sets by our algorithm to *completely* cover the terrain, we get only an $O(\log n)$ approximation. The same applies if we compare our solution to an optimal ε-cover. However, in our approach we simultaneously approximate the amount of terrain covered, and the number of guards required. This allows us to find a decent approximation ratio. We further remark that the minimum number of guards required for an ε-cover may differ significantly from the minimum number of guards to completely cover the terrain.

Finally, in case we measure the size of a viewshed by its area we can compute an ε-cover of $\mathcal{T}'$ in $O(n^2(m \log n + m^4 + k^4/\varepsilon^4))$ time using the results of Hurtado *et al.* [14]. The $O(n^2m^4)$ term originates from computing $[\![\mathcal{V}(\mathcal{P})]\!] = [\![\mathcal{T}']\!]$. If we know the size of $\mathcal{T}'$ in advance, for example because $\mathcal{T}' = \mathcal{T}$, then we can improve these results to $O(n^2(m \log n + k^4/\varepsilon^4))$.

terrain	Hot Springs	Quinn Peak	Sphinx Lakes	Split Mountain	Wren Peak
	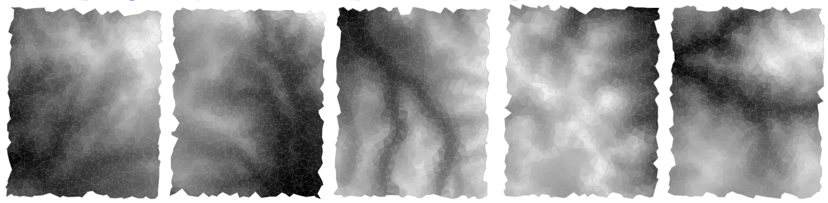				
coarse	1797	1787	1778	1769	1781
fine	16 197	16 087	16 017	15 895	16 065

Table 1. The terrains we use to evaluate our approach. The terrains are shaded by height: higher vertices and faces are lighter.

Experimental Evaluation

We experimentally evaluate results from the greedy algorithm on five terrain models in California, obtained from the U.S. Geological Survey[3] and shown in Table 1. Each terrain model spans approximately 11.5km $\times$ 14km. For each terrain we have a coarse and a fine version with approximately 1800 and 16 000 vertices, respectively. In all our experiments we choose the set $\mathcal{P}$ of potential guards to be the set of vertices of the terrain.

To keep the implementation of our algorithms simple, we consider only the case where the size of the viewsheds is measured by the number of terrain vertices. For the same reason we use the naive implementation for the point-to-point visibility rather than building the data structure for visibility queries.

We first investigate the number of guards selected, that is, the size of the set $\mathcal{G}$, as a function of ε. For these results we fix the height h of the guards on 15 meters. The results are shown in Fig. 2 and 3. For both the coarse and fine terrain models we see that the required number of guards rapidly increases when $1-\varepsilon$ approaches 1, that is, when we attempt to cover almost the complete

[3] http://www.usgs.gov

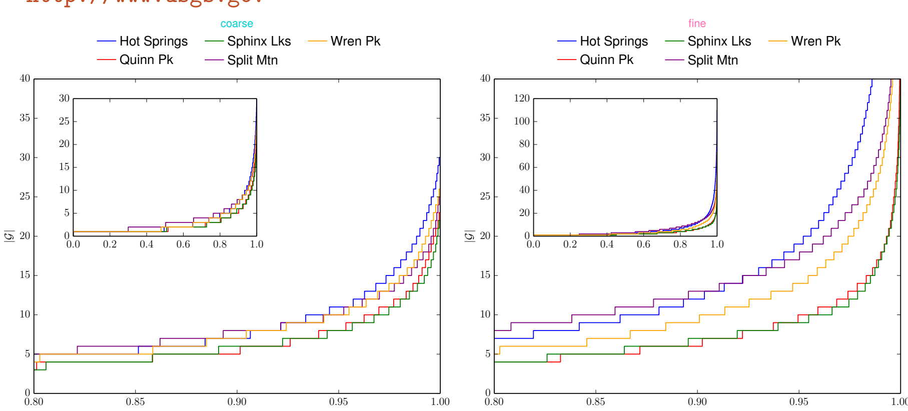

Fig. 2. The number of selected guards in $\mathcal{G}$ as a function of ε. The insets show the results for full domain of ε.

terrain. However, eleven guards are sufficient for an 0.05-cover on each coarse terrain model. Generally speaking, the same number of guards covers a smaller fraction of the terrain in the fine terrain models than in the coarse terrain models, indicating that the added detail may be significant for visibility studies. For the coarse terrain models eleven guards can cover at least 88% of the terrain; for an 0.05-cover we need between nine and nineteen guards, depending on the terrain.

Fig. 4 shows the locations of the guards placed by GreedyGuard for an 0.05-cover on the coarse Wren Peak terrain (again for height $h = 15\text{m}$). The fraction of the terrain covered by those guards is shown in Fig. 3.

The height of the guards influences the number of guards required, see Fig. 5. However, the differences are small. The height has a larger influence on the fine version of the terrains. Even so, covering one of the terrains with guards at height 1m requires only four more guards than guards at 30m.

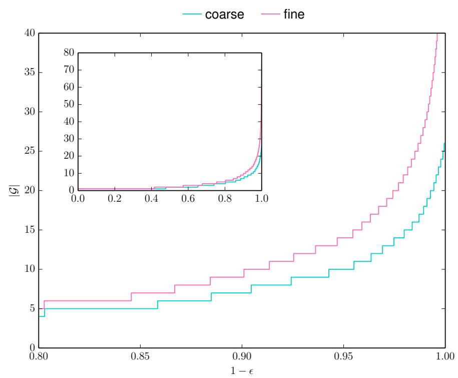

Fig. 3. A comparison between the number of guards needed in the coarse and fine versions of the Wren Peak terrain.

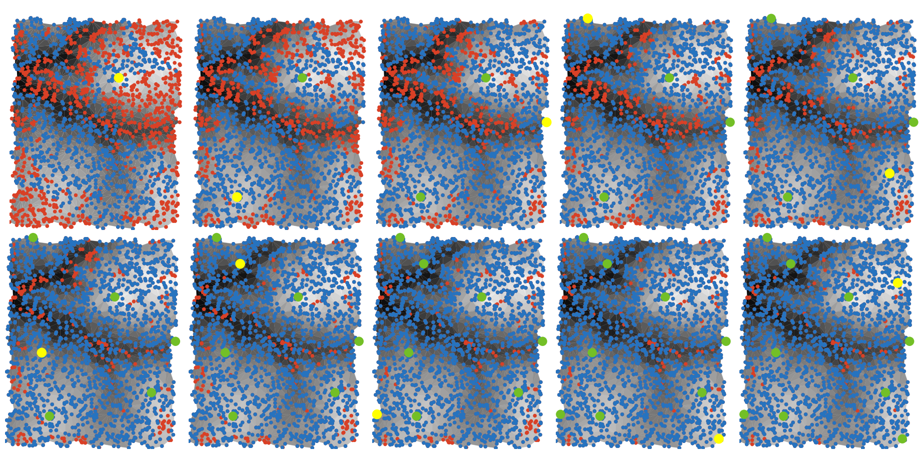

Fig. 4. The ten guards placed by GreedyGuard for a 0.05-cover of the coarse Wren Peak terrain (green). Each figure shows the vertices covered so far in blue, and the vertices that remain uncovered in red. The newly selected guards are shown in yellow.

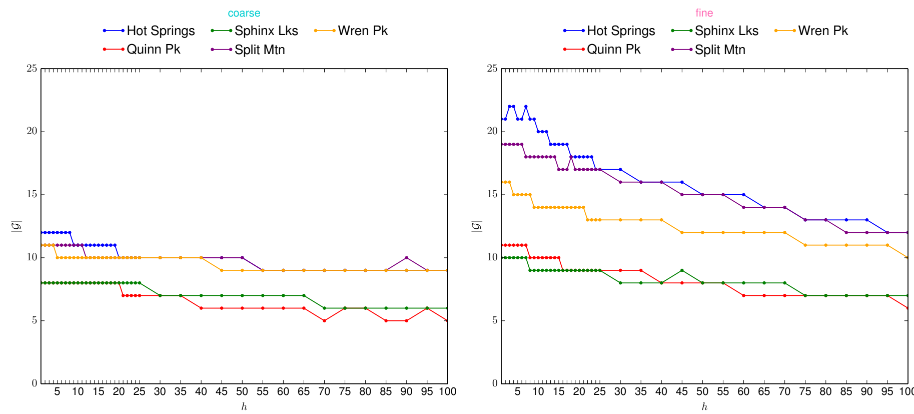

Fig. 5. The number of required guards for an 0.05-cover as a function of the heights h.

4 Dominating Guards

4.1 Simple Domination

A guard g *dominates* another guard h if (and only if) the viewshed of h is contained in the viewshed of g, that is, $\mathcal{V}(h) \subseteq \mathcal{V}(g)$. We say g *strictly* dominates h if $\mathcal{V}(h) \subset \mathcal{V}(g)$. We now observe:

Observation 1. *Let $\mathcal{P}$ be a set of potential guards. There is an optimal (minimum size) ε-cover $\mathcal{G}$ of $\mathcal{V}(\mathcal{P})$ such that no guard in $\mathcal{G}$ is strictly covered by any guard in $\mathcal{P}$.*

Proof. Let $\mathcal{G}^*$ be an optimal ε-cover of $\mathcal{V}(\mathcal{P})$. Let h be any guard $h \in \mathcal{G}^*$ that is strictly dominated by another potential guard $g \in \mathcal{P}$. Replace h by g. Repeat this procedure until there are no more guards that satisfy the criterion. Note that this process terminates since strict domination defines a partial order on $\mathcal{P}$ (a guard g can only strictly dominate a guard h if $[\![\mathcal{V}(g)]\!] > [\![\mathcal{V}(h)]\!]$). Let $\mathcal{G}$ be the set of guards we obtain. Clearly, $\mathcal{G}$ contains at most the same number of guards as $\mathcal{G}^*$, and since we have that $\mathcal{V}(\mathcal{G}^*) \subseteq \mathcal{V}(\mathcal{G})$, $\mathcal{G}$ is also an ε-cover of $\mathcal{V}(\mathcal{P})$. $\square$

It now follows that we never have to choose a guard that is strictly dominated by another guard. Hence, when finding an ε-cover of the terrain, we can simply ignore all dominated guards. Although there is no useful theoretical lower bound on the number of dominated guards, we will see that for the terrains considered here, we can discard between 10% and 40% of the potential guards, depending on the resolution of the terrain.

Experimental Evaluation

In Fig. 6 shows the percentage of the potential guards that is dominated as a function of the height h, and Fig. 7 shows the dominating and dominated potential guards for $h = 15\text{m}$. For the coarse terrains, between forty and fifty percent of the potential guards are dominated by another potential guard, and can thus be discarded. For the fine terrains the percentage of dominated vertices is lower: between eight and twenty percent. Surprisingly, the percentage of dominated vertices in the fine Hot Springs terrain is well over twenty percent. This behavior differs from the coarse Hot Springs terrain, where the percentage of dominated guards is comparable to the other terrains.

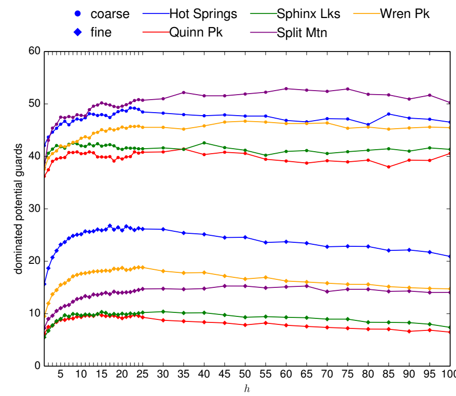

Fig. 6. The percentage of the potential guards that are dominated by another guard as a function of their height.

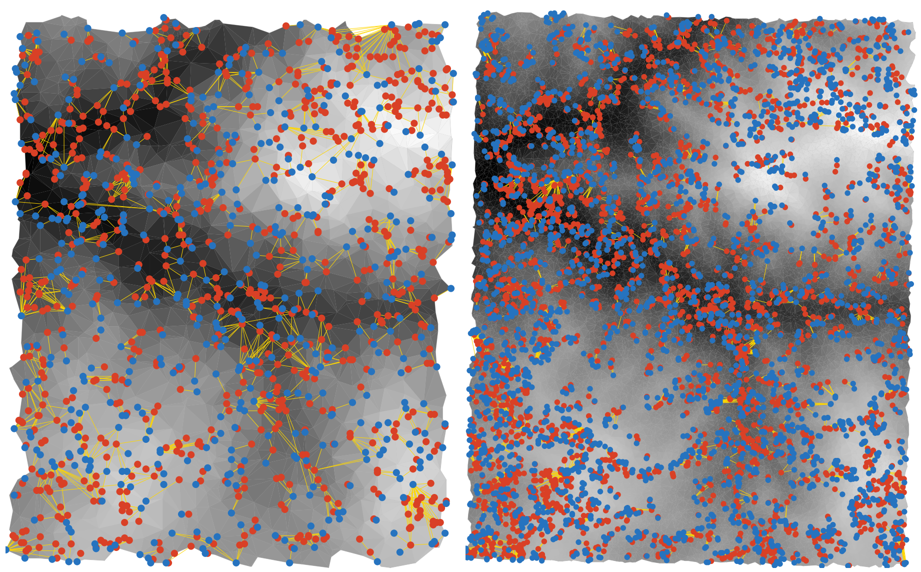

Fig. 7. Dominating potential guards in blue and the potential guards they dominate in red on the Wren Peak terrain. On the left the coarse model and on the right the fine model.

The height at which we place the guards only mildly influences the number of dominated vertices. One might expect that the number of dominated vertices increases monotonically as the height grows (since the individual viewsheds grow when the height increases). However, for the heights considered this is not the case: the number of dominated vertices increases at first, but then slightly decreases.

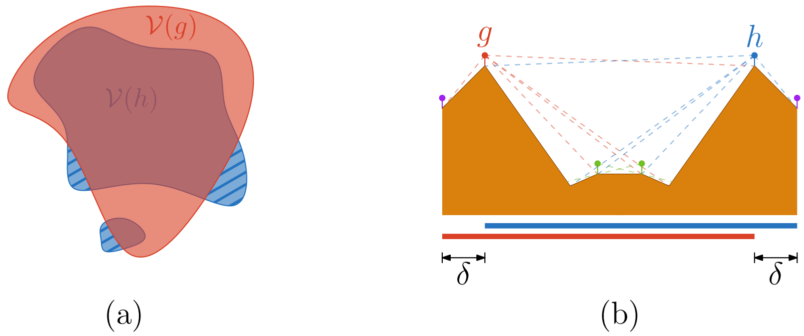

Fig. 8. (a) The red guard g δ-dominates the blue guard h if and only if the fraction of $\mathcal{V}(h)$ that is not covered by g (the hatched area) is at most δ. (b) A cross section of a terrain in which all potential guards are δ-dominated by an other guard.

4.2 δ-Domination

Since it is sufficient if the guards cover a fraction of $1 - \varepsilon$ of the terrain, we can consider slightly relaxing the definition of domination. Instead of requiring that guard g should see everything that h sees, it may be sufficient if g sees a large enough portion of h. More formally, guard g δ-*dominates* guard h if (and only if) the fraction of $\mathcal{V}(h)$ that is not covered by g is at most δ, that is, if and only if $[\![\mathcal{V}(h) \setminus \mathcal{V}(g)]\!] \,/\, [\![\mathcal{V}(h)]\!] \leq \delta$. See Fig. 8(a) and 9.

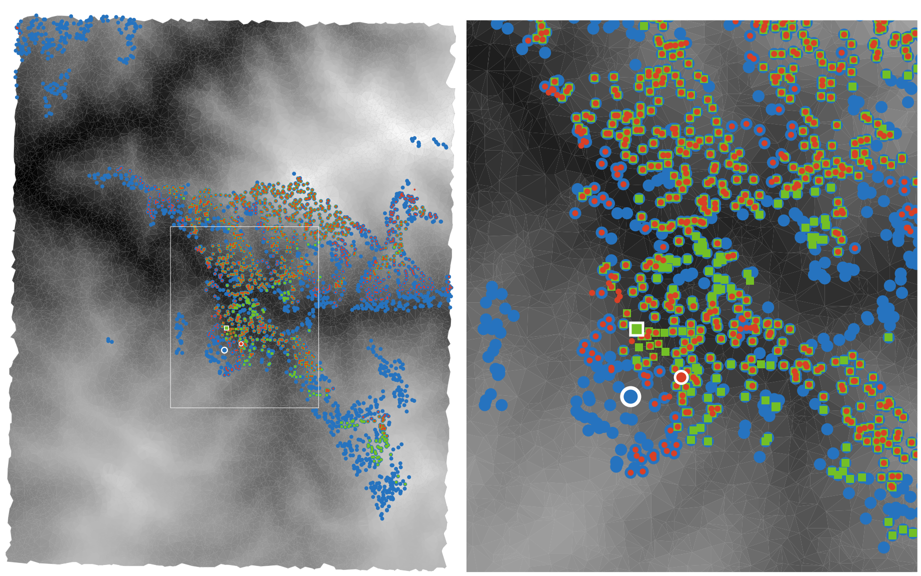

Fig. 9. Three guards on a fine grained Wren Peak terrain (left). The blue guard 0.02-dominates the red and green guards (guards are delineated in white). Only very few vertices in the red and green viewsheds are not visible from the blue guard (right).

In this setting we can no longer just discard all guards that are δ-dominated, as shown in Fig. 8(b). For small values of ε, the set $\mathcal{G} = \{g, h\}$ is the only optimal solution. However, guard g δ-dominates h, and vice versa. Thus, we cannot remove them both. Instead, we extend the notion of δ-domination to sets of guards: a set of guards $\mathcal{G}$ δ-dominates a set of guards $\mathcal{H}$ if for every guard in $\mathcal{H}$ there is a guard in $\mathcal{G}$ that δ-dominates it. We now wish to find a minimum size set of guards that δ-dominate $\mathcal{P}$.

It turns out that computing such a minimum size set is NP-hard as well [13]. However, we can easily find a *maximal* set $\mathcal{D}$ with a simple greedy algorithm. Unfortunately, we cannot provide a good lower bound on the fraction of the terrain that is visible from $\mathcal{D}$, so instead we evaluate δ-domination experimentally.

Experimental Evaluation

In Fig. 10 we can see the percentage of the guards for which there is another guard that δ-dominates it (for $h = 15$ meters). This figure shows that on the fine terrains and $\delta \geq 0.05$ for practically every vertex there is another vertex that δ-dominates it. For the coarse terrains we need slightly higher values of δ. This is to be expected, since the vertices, and thus the locations for the potential guards, are spread further apart than in the coarse terrains. Even so, we note that for $\delta = 0.05$ already 90% of the potential guards are δ-dominated by another guard.

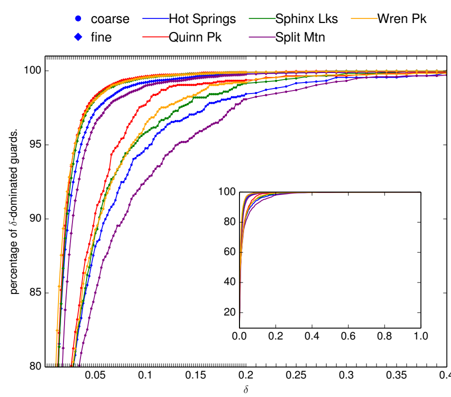

Fig. 10. The percentage of potential guards in $\mathcal{P}$ for which there is an other potential guard that δ-dominates it.

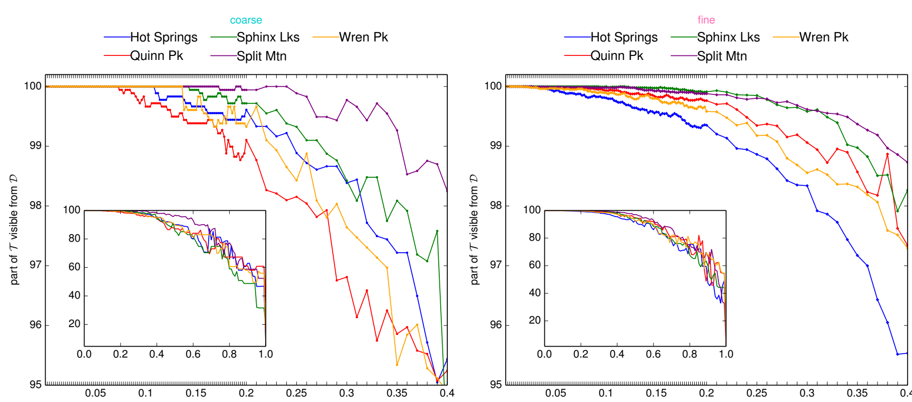

Fig. 11. The fraction of $\mathcal{T}$ that it still coverable with a minimal set $\mathcal{D}$ of δ-dominating guards.

As noted before, we can no longer remove *all* guards that are δ-dominated, so instead consider a minimal size set $\mathcal{D}$ of δ-dominating guards. Fig. 11 shows the percentage of the terrain that is still coverable by $\mathcal{D}$ (again for $h = 15$ meters). For values of δ up to 0.2—so 20% of the viewshed of a guard h may not be visible from a guard g that δ-dominates it—a minimal set $\mathcal{D}$ of δ-dominating guards can still see more than 99% of the terrain.

5 The Greedy Approach with δ-Domination

We can now use the δ-domination to limit the number of potential guards we have to consider when computing an ε-cover. As a preprocessing step we determine which guards are δ-dominated, so we can discard some of the potential guards. We then use the greedy algorithm from Section 3 to compute an ε-cover on the remaining guards. More specifically, consider the following algorithm DominatingGuard, that computes an ε-cover using the δ-domination, if possible. When the fraction of the terrain $\hat{\delta}$ that is no longer visible from $\mathcal{D}$ is larger than ε we (obviously) cannot obtain an ε-cover any more.

Algorithm DominatingGuard$(\mathcal{T}, \varepsilon, \delta, \mathcal{P})$
1. Compute the viewsheds for all guards in $\mathcal{P}$.
2. Compute a minimal set of guards $\mathcal{D}$ that δ-dominates $\mathcal{P}$.
3. Let $\hat{\delta} = [\![\mathcal{V}(\mathcal{D})]\!] \,/\, [\![\mathcal{V}(\mathcal{P})]\!]$ be the fraction of $\mathcal{V}(\mathcal{P})$ covered by $\mathcal{D}$.
4. Let $\gamma = (\varepsilon - \delta)/(1 - \hat{\delta})$ and let $\hat{\mathcal{T}} = \mathcal{V}(\mathcal{D})$.
5. **return** GreedyGuard$(\hat{\mathcal{T}}, \gamma, \mathcal{D})$

We note, however, that the terrain can be preprocessed by removing δ-dominated vertices, irrespective of the algorithm used to then place the guards. Since the data size can be reduced drastically, it may also be feasible to use more computation-intensive solutions than the simple greedy heuristic presented in this study.

Experimental Evaluation

We choose $\varepsilon = 0.05$ and $h = 15$m and run DominatingGuard for various values of δ. Fig. 12 shows the number of guards in the set $\mathcal{D}$, as a function of δ. The number of potential guards to be given to the greedy algorithm rapidly decreases when δ increases. On the fine grained terrains and for $\delta \geq 0.04$ we can discard well over 90% of all potential guard locations due to the δ-domination. For the coarse terrains this number is roughly 80%.

We observed that the number of guards selected is independent of δ: for all values the algorithm returns the same number of guards. Hence, removing δ-dominated potential guards does not seem to affect the final solution.

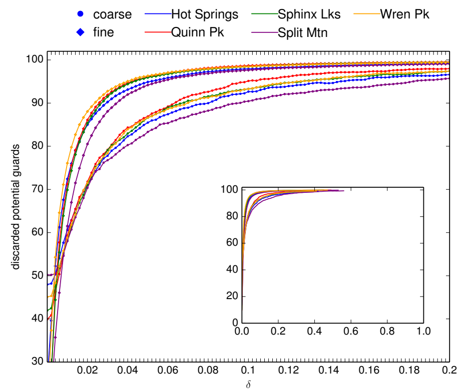

Fig. 12. The percentage of guards that we can discard in step two of DOMINATING-GUARD. When the selected minimal set of dominating guards no longer cover at least $1 - \varepsilon$ of the terrain the algorithm does not yield a solution.

6 Concluding Remarks

We investigated practical approaches to compute an ε-cover of a polyhedral terrain $\mathcal{T}$; that is, a set of guards that together can see at least a fraction $1 - \varepsilon$ of the terrain. We showed that, for any ε of interest, no constant-approximation algorithm exists to compute an ε-cover. In addition, we provided a theoretical analysis of a straightforward greedy algorithm to compute an ε-cover. This analysis shows that we need at most $O(k/\varepsilon)$ guards, where k is the minimum number of guards required to completely cover the terrain. Through experiments we show the algorithm gives a reasonable number of guards in practice. Furthermore, we introduced the notions of dominated and δ-dominated guards. Our experiments show that these can greatly reduce the number of potential guard locations we have to consider.

In our experiments we use viewsheds that only consider the vertices of the terrain. It would be very interesting to see if we obtain equally good numbers for viewsheds that also incorporate the (interior of) the faces of the terrain. A different direction of future work entails a theoretical investigation of the concept of δ-domination. For our current definition, we cannot state any theoretical guarantees, though it performs well in practice.

Acknowledgments.

This work has been partially supported by the Netherlands Organisation for Scientific Research (NWO) under grants 639.021.123 and 612.001.022.

References

1. Agarwal, P.K., Har-Peled, S., Varadarajan, K.R.: Geometric approximation via coresets. In: Combinatorial and Computational Geometry, MSRI, University Press (2005) 1–30
2. Alon, N., Moshkovitz, D., Safra, S.: Algorithmic construction of sets for k-restrictions. ACM Transactions on Algorithms (2006) 153–177
3. Aspbury, A.S., Gibson, R.M.: Long-range visibility of greater sage grouse leks: a gis-based analysis. Animal Behaviour **67**(6) (2004) 1127–1132
4. Burrough, P.A., McDonnell, R., Burrough, P.A., McDonnell, R.: Principles of geographical information systems. Volume 333. Oxford university press Oxford (1998)
5. Camp, R.J., Sinton, D.T., Knight, R.L.: Viewsheds: A complementary management approach to buffer zones. Wildlife Society Bulletin **25**(3) (1997) 612–615
6. Catry, F.X., Rego, F.C., Santos, T., Almeida, J., Relvas, P.: Fires prevention in Portugal - using GIS to help improving early fire detection effectiveness. In: 4th Int. Wildland Fire Conf. (2007)
7. Cole, R., Sharir, M.: Visibility problems for polyhedral terrains. Journal of Symbolic Computation **7**(1) (1989) 11–30
8. De Floriani, L., Marzano, P., Puppo, E.: Line-of-sight communication on terrain models. International journal of geographical information systems **8**(4) (1994) 329–342
9. Downey, R.G., Fellows, M.R.: Parameterized Complexity. Springer-Verlag (1999) 530 pp.
10. Eidenbenz, S.J., Stamm, C., Widmayer, P.: Positioning guards at fixed height above a terrain – an optimum inapproximability result. In: Algorithms ESA 98. Volume 1461 of Lecture Notes in Computer Science. Springer (1998) 1–12
11. Floriani, L.D., Magillo, P.: Algorithms for visibility computation on terrains: a survey. Environ. Plann. B **30**(5) (2003) 709–728
12. Franklin, W.: Siting observers on terrain. In Richardson, D., Oosterom, P., eds.: Advances in Spatial Data Handling. Springer Berlin Heidelberg (2002) 109–120
13. Garey, M.R., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-completeness. WH Freeman & Co. New York, NY, USA (1979)
14. Hurtado, F., Löffler, M., Matos, I., Sacristán, V., Saumell, M., Silveira, R.I., Staals, F.: Terrain visibility with multiple viewpoints. In: Proc. 24th International Symposium on Algorithms and Computation. (2013)
15. Johnson, D.S.: Approximation algorithms for combinatorial problems. Journal of computer and system sciences **9**(3) (1974) 256–278
16. Jones, N.L., Wright, S.G., Maidment, D.R.: Watershed delineation with triangle-based terrain models. Journal of Hydraulic Engineering **116**(10) (1990) 1232–1251
17. Katz, M.J., Overmars, M.H., Sharir, M.: Efficient hidden surface removal for objects with small union size. Computational Geometry **2**(4) (1992) 223–234
18. Kim, Y.h., Rana, S., Wise, S.: Exploring multiple viewshed analysis using terrain features and optimisation techniques. Computers & Geosciences **30** (2004) 1019–1032
19. Kreveld, M.: Digital elevation models and tin algorithms. In Kreveld, M., Nievergelt, J., Roos, T., Widmayer, P., eds.: Algorithmic Foundations of Geographic Information Systems. Volume 1340 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (1997) 37–78

20. Lv, P., Zhang, J., Lu, M.: An optimal method for multiple observers sitting on terrain based on improved simulated annealing techniques. In: Advances in Applied Artificial Intelligence. Volume 4031 of LNCS. Springer (2006) 373–382
21. Magalhaes, S.V.G., Andrade, M.V.A., Franklin, W.R.: An optimization heuristic for siting observers in huge terrains stored in external memory. In: Hybrid Intelligent Systems (HIS), 2010 10th International Conference on. (Aug 2010) 135–140
22. Silfer, A.T., Kinn, G.J., Hassett, J.M.: A geographic information system utilizing the triangulated irregular network as a basis for hydrologic modeling. In: Proc. Auto-Carto. Volume 8. (1987) 129–136
23. Silveira, R.I.: Optimization of polyhedral terrains. ASCI Disseratation Series (178) (2009)
24. Stephen, B.K., Wicker, S.B.: Experimental analysis of local search algorithms for optimal base station location (2000)