Journal of Interconnection Networks

STATION LAYOUTS IN THE PRESENCE OF LOCATION CONSTRAINTS

PROSENJIT BOSE and EVANGELOS KRANAKIS
Carleton University, School of Computer Science
Ottawa, Ontario, K1S 5B6, Canada

CHRISTOS KAKLAMANIS and LEFTERIS M. KIROUSIS
University of Patras, Department of Computer Engineering and Informatics
GR-265 04 Patras, Greece

DANNY KRIZANC
Wesleyan University, Department of Mathematics
Middletown CT 06459, USA.

DAVID PELEG
Weizmann Institute of Science
Department of Computer Science and Applied Mathematics
Rehovot 76100, Israel

Received (received date)
Revised (revised date)

In wireless communication, the signal of a typical broadcast station is transmitted from a broadcast center p and reaches objects at a distance, say, r from it. In addition there is a radius r_0, $r_0 < r$, such that the signal originating from the center of the station is so strong that human habitation within distance r_0 from the center p should be avoided. In other words, points within distance r_0 from the station comprise a hazardous zone. We consider the following *station layout* problem: Cover a given planar region that includes a collection of buildings with a minimum number of stations so that every point in the region is within the reach of a station, while at the same time no interior point of any building is within the hazardous zone of a station. We give algorithms for computing such station layouts in both the one- and two-dimensional cases.

Keywords: Approximation algorithm, Broadcast Station, Health hazards, Optimal Layout, Wireless Communication.

1. Introduction

In wireless communication we are interested in providing access to communication to a region (e.g., a city, a campus, etc.) within which several sites (e.g., buildings) are located. Closeness to stations may be undesirable in certain instances, e.g., hospital or laboratory facilities, people with heart pace-makers, etc. Characteristically, closeness to broadcasting stations is undesirable in the case of mobile telephony. Cellular phones are radio receivers which operate in the ultra-high frequency (UHF) band. They receive radio transmissions from a central base station (or cell) at frequencies between 869 and 894 MHz and retransmit their radio signal back to the base station at frequencies between 824 and 850 MHz. Continuous exposure to strong radio transmissions is considered to be hazardous See, for example, Lin[8] for a comprehensive study and survey of the biological effects of exposure to radio frequency resulting from the use of mobile and other personal communication services. The strength of a radio signal is inversely proportional to the square of the distance from the emitting station. Therefore, sites close to base stations of mobile telephony are continuously exposed to hazardous transmissions. Nevertheless, there is a threshold in a signal's strength ($1W$ is the currently accepted value) below which the signal is sufficiently safe but still strong enough to reach its destination. This makes it desirable to build cellular telephony base stations far from existing buildings, but densely enough to cover a given area.

As we show in this paper, the problem of station layout so that certain locations are not close to any station can be reduced to the problem of covering a given area with a minimum number of stations placed on points that are selected from a given finite collection of permissible locations. Relevant to this problem is the recent installation in several major U.S. cities of Ricochet network transceivers that provide wireless access to data networks. These are radio transceivers mounted on existing streetlight or utility poles[10]. The selection of the poles has already raised concern about possible health hazards (see Horowitz[7]).

In this paper we consider broadcast station layouts in wireless communication that take into account health hazards resulting from the closeness of human habitation to the transmission station. Given such constraints we are interested in minimizing the number of broadcast stations used. The buildings are located within a region R. For the sake of simplicity we assume R to be rectangular. In the most general case the buildings may be represented by simple polygons with or without holes.

1.1. *Formulation of the problem and notation*

The parameters involved in transmissions for a typical station in the plane are the transmission center p of the station, and positive real numbers $r_0 < r$ such that

- r determines the range of the station, i.e., the signal transmitted from the center p can reach any destination at distance r from the center.
- r_0 determines the hazardous zone of the station, i.e., the strength of the transmitted signal exceeds permissible health constraints within distance r_0 from the center.

Throughout this paper, we present our results in terms of the L_1-distance (also called Manhattan metric). So, the square $D(p; r_0) = \{x : d(x, p) \leq r_0\}$ is the locus

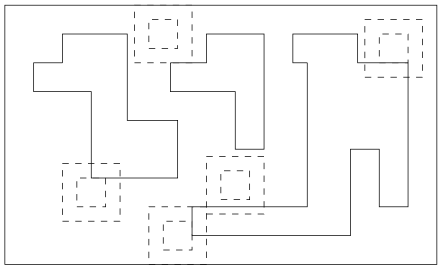

Figure 1: A rectangular region with polygonal buildings to be covered by stations so that no hazardous zone of a station intersects the interior of a building.

of points that are "too close" to the broadcast center p. Existing health constraints make it advisable that human habitation is not allowed within the square $D(p; r_0)$. On the other hand, the square $D(p; r) = \{x : d(x, p) \leq r\}$ is the locus of points reachable from the station centered at p. The values of the parameters r_0, r are suggested by the manufacturer.

Consider a rectangular area R and a set $\mathcal{P}$ of polygons (the buildings) that lie inside R (the polygons may have holes). We assume that all edges of the buildings are parallel to one of the Cartesian axes. We examine the following problem:

Problem 1 (Station Layout Problem — SLP) *Either find a set of points $\{p_1, \ldots, p_m\} \subseteq R$ of minimum cardinality such that (i) the set of squares $D(p_i; r), i = 1, \ldots, m$ cover R and (ii) no point in the interior of any of the polygons in $\mathcal{P}$ lies in any of the squares $D(p_i; r_0), i = 1, \ldots, m$, or report that no such set of points exists.*

Notice the distinction between boundary and interior points to a building. Points in the boundary of the buildings are allowed to lie within a hazardous zone, but interior points are not. In addition, a point may lie within the range of more than one station. Figure 1 depicts a rectangular region with buildings (drawn with solid lines) and stations whose range is depicted by squares with dashed lines.

We use the term *unit square centered at a point p* to refer to the square $D(p; 1/2)$. Observe that SLP is computationally equivalent to the following problem:

Problem 2 (Reduced Station Layout Problem — red-SLP) *Either find a set of points $\{p_1, \ldots, p_m\} \subseteq R$ of minimum cardinality such that (i) the set of unit squares centered at $p_1, \ldots, p_m$ cover R and (ii) none of the points $p_1, \ldots, p_m$ lies in the interior of any of the polygons in $\mathcal{P}$, or report that no such set of points exists.*

Clearly, SLP is more general than red-SLP. However, the former can also be reduced to the latter. To provide this reduction, we surround each building with a strip of width r_0 and merge the resulting polygons into the new buildings. We also change units so that $r = 1/2$.

In the sequel, the term *cover* of the area R will always refer to a cover of R with unit squares whose centers do not lie in the interior of any building in $\mathcal{P}$.

Of course, by a change in units, we can assume that the covering squares of an instance of red-SLP have edge-length not 1, but an arbitrary given number.

Notice the equivalance of intances of Problem 2 with the Euclidean metric (covering elements are unit circles) and Manhattan metric (covering elements are unit squares). To see this observe that a unit circle is included inside a square of side 2, and a unit square is included inside a circle of radius $\sqrt{2}$. Also the assumption that the edges of the polygons are axis-parallel is not essential to the theory, because an arbitrary polygon can be approximated to any desirable degree by an axis-parallel polygon by placing it on a grid of sufficiently large resolution. We leave to the reader the necessary additional details of how to prove rigorously the equivalence of these problems to our original problem.

1.2. *Results of the paper*

In the next section, we consider the one-dimensional analog of the problem. We deal with the two-dimensional case in Section . In the latter section, we first consider an algorithm for testing the existence of a solution. Subsequently, we show how to reduce the problem to a discrete problem in which the centers of the stations are to be selected from predetermined points within the region R. This is used to provide:

- a simple, polynomial time, logarithmic approximation algorithm and
- a polynomial time, constant approximation, dynamic programming algorithm

for red-SLP and hence for SLP.

Finally in Section we examine the case of "thin" buildings and we give for such buildings a linear time, constant approximation algorithm for SLP.

1.3. *On the number of stations*

Notice that even if we keep the total area of the buildings bounded, the least number of unit squares required to solve instances of red-SLP may be unbounded. Indeed, let R be the the square area of edge-length 1.5 depicted as $ABCD$ in Figure 2a and let $\mathcal{P}$ contain a single polygon, the one delimited by the line segments AE, FC, CG, HA and the two step-lines EF and GH (each with 13 turn-points). Notice that to cover the points of the interior of the polygon that belong to the middle square $KLMN$ of edge-length 1/2, we need at least 6 unit squares with centers on the boundary of $\mathcal{P}$. For example, centering six unit squares on the vertices p_1, p_2, p_3 of the step-line EF and p_4, p_5, p_6 of the step-line GH, respectively, $KLMN$ is covered (see Figure 2b). This cannot be accomplished with less than 6 unit squares. With two more unit squares centered at p_7 and p_8, respectively, the whole area R is covered. If each of the step-lines has $2n+2$ turn-points instead of 13, $n+2$ unit squares are needed to cover R.

As a consequence we get that as input size to the problem we should take

$$N = \text{Area}(R) + n,$$

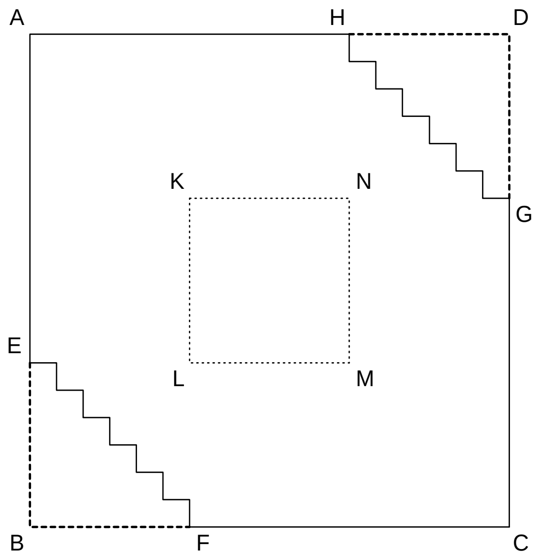

(a)

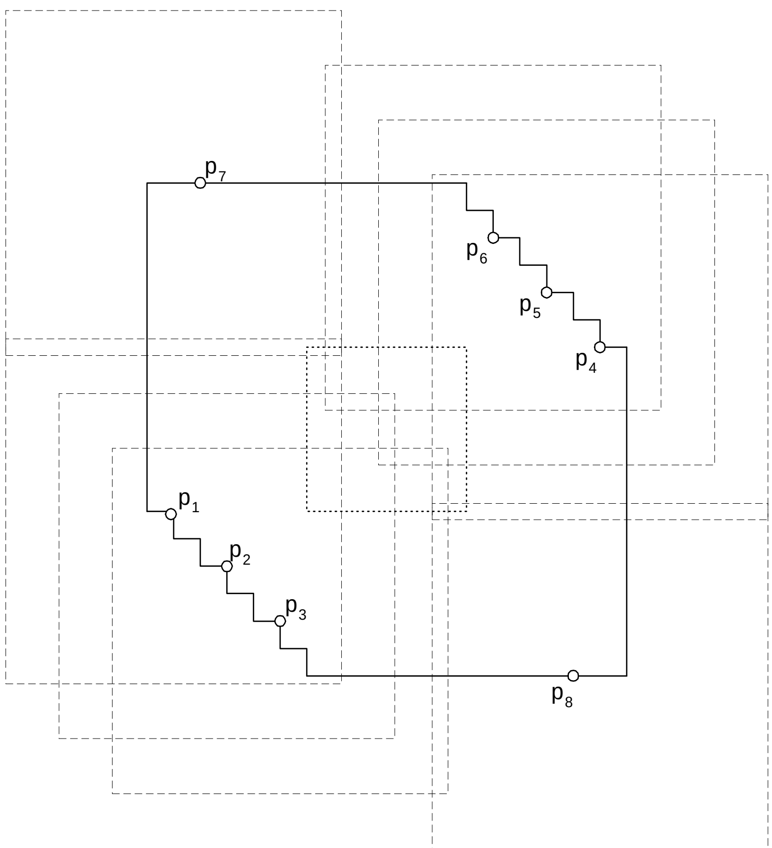

(b)

Figure 2: A square region of edge-length 1.5 and a polygon in it so that the number of unit squares with centers not in the interior of the polygon that are needed to form a cover is $\Theta(n)$, where n is the number of vertices of the polygon. The unit squares have been drawn with edge-length slightly larger than 1, to avoid coincidences that blur the picture.

where n is the number of vertices of all the polygonal buildings. Notice that Area(R) is a straightforward lower bound on the size of the solution and that, even for constant Area(R), $\Omega(n)$ stations might be required, as the example shows.

2. Algorithm on the Line

In this section, we consider the one-dimensional analog of the Reduced Station Layout Problem (1-red-SLP). In this case, the area R is an interval, say I. Unit "squares" are intervals of length 1. Also intervals, say $I_i, i = 1, \dots, n$, are the "buildings" in $\mathcal{P}$. Merging intersecting intervals, we assume that the "buildings" are actually pairwise disjoint. By sorting their endpoinds we can also assume that the sequence of endpoints is $o_1 < o_2 < o_3 < o_4 < \cdots < o_{2n-1} < o_{2n}$, where o_i is the left (respectively, right) endpoint of an interval for i odd (respectively, even). Let ℓ be the maximum length of a building, i.e., $\ell = \max\{\ell(I_i) : i = 1, \dots, n\}$, where $\ell(I_i) = o_{2i} - o_{2i-1}$, for $i \geq 2$.

Lemma 1 *An instance of 1-red-SLP has a solution if and only if $\ell \leq 1$. Hence there exists a linear time algorithm for determining solvability of 1-red-SLP.*

PROOF. . The condition is obviously necessary. To prove the sufficiency, we cover the set of buildings by placing a station at each endpoint of each interval, as well as at remaining gaps between buildings. ■

Now we prove the main theorem of this section.

Theorem 1 *There exists an $O(N \log N)$ time algorithm for computing a minimum size cover for an instance of 1-red-SLP.*

PROOF. Given the previous lemma, we may restrict our attention to the case when a solution exists. The problem can be solved by a greedy algorithm for locating the centers of the stations. We scan the given intervals "moving" from left-to-right and we place the stations successively as far to the right as possible without leaving any (part of a) building uncovered. To be more precise, the algorithm proceeds as follows: First assume that there is an interval endpoint lying to the right of the rightmost already covered point and within distance $\leq 1/2$ from it. In this case, consider the rightmost such endpoint. If this endpoint is odd numbered (i.e., it is the left endpoint of an interval) then place the midpoint of the station there. If this endpoint is even numbered (i.e., it is the right endpoint of an interval), then if there is an overlap of size $\epsilon \geq 0$ between the area previously covered and the range of the station that is now to be placed (i.e., if the station were to be placed at the right endpoint), place the midpoint of a station ϵ units to the right. Here it should be observed that this will not place a midpoint of a station inside a building because of the choice of the rightmost endpoint before. If on the other hand there is no interval endpoint lying to the right of the rightmost already covered point and within distance $\leq 1/2$ from it, then center the next station at distance $1/2$ to the right of the rightmost already covered point. The algorithm terminates when the entire interval I is covered by the stations.

To prove minimality observe that we can show by induction that for every point

x on I, we find a placement that covers x and everything to the left of x with the minumum number of stations, and our solution also covers as much to the right of x as possible using this many stations.

The $O(N \log N)$ bound arises from sorting interval endpoints required for preprocessing. This completes the proof of the Theorem. ∎

3. Algorithms on the Plane

In this Section, we give approximation algorithms for the plane version of red-SLP (and hence SLP).

3.1. *Testing for a solution*

Theorem 2 *A solution for red-SLP exists if and only if there is no unit square that lies entirely in the interior of any polygon. In particular, there exists an $O(n^2)$ time algorithm to determine whether or not a solution exists, where n is the number of vertices in all the polygons (i.e., buildings).*

PROOF. The condition is clearly necessary. To prove that it is sufficient we argue as follows. Take any building in $\mathcal{P}$ and cover it with unit squares in the following greedy manner. Move along the building's perimeter by placing unit squares with centers on the perimeter. For each edge of the building start by placing a unit square centered at one endpoint and place centers of unit squares on this edge at a distance 1 apart, and ending with a unit square centered at the other endpoint. In view of the condition above, the polygon will be fully covered. The remaining area of the rectangular region can of course always be covered with unit squares. From an algorithmic point of view, it forms a simple polygon with holes and can be covered with a standard (approximation) algorithm for square covering, e.g., Aupperle[1], and Bar-Yehuda[2]. This algorithm gives a solution, if there is one, however it gives no indication of the optimality of the number of stations used.

Nevertheless, we have proved that to test if a solution exists, it is enough to show that no unit square lies in the interior of a polygon. This can be easily tested with an $O(n^2)$ time algorithm for finding a maximum square inside a polygon (see De Pano *et al* [5]). This completes the proof of the Theorem. ∎

3.2. *Reduction to a discrete problem without buildings*

Now we reduce red-SLP to a discrete problem without buildings. Given an instance of red-SLP determined by a rectangular region R and a collection of polygons $\mathcal{P}$, define a collection $\mathcal{Z}$ of points in R as the union of the two sets $\mathcal{Z}_0$ and $\mathcal{Z}_1$ defined below:

- Partition the rectangular region R into parallel and horizontal strips at distance $1/2$ apart and let $\mathcal{Z}_0$ be the collection of points of intersection of these strips which lie neither in the interior nor on the boundary of any building in $\mathcal{P}$ (the last horizontal and the last vertical strip may be of length $\leq 1/2$).

- The points in $\mathcal{Z}_1$ lie on the perimeter of the buildings in $\mathcal{P}$. These points are of two types: (a) all vertices of these polygons, (b) for any polygon in $\mathcal{P}$, and starting from an arbitrary vertex of the polygon, walk along the perimeter and place points on the perimeter at distance 1/2 apart (the walk terminates when we return to the starting point).

We refer to unit squares whose centers are points in $\mathcal{Z}$ as *$\mathcal{Z}$-squares*. A *discrete* cover of the region R is a cover by $\mathcal{Z}$-squares. Since an edge of a building, being a line segment in R, may have length at most Area(R) (e.g., when R is "thin"), the number of points in $\mathcal{Z}_1$ is at most equal to the product of the number of the vertices of the buildings with Area(R). Therefore, the number of $\mathcal{Z}$-squares is at most N^2. Indeed, we know that $|\mathcal{Z}_0| \leq$ Area(R), and that for any individual polygon, p_i, the number of points needed to cover it is at most $n_i + \lceil(\text{perimeter}(p_i)\rceil - 1$ (where n_i is the number of vertices in p_i). The -1 term comes because the first point laid out in a walk of the polygon is also a vertex. The sum of the n_i is n, and the sum of the ceilings of the perimeters is certainly at most $n \cdot \text{Area(R)} + n$, so we have $|\mathcal{Z}_1| \leq n \cdot (\text{Area(R)} + 1)$. Now since $N = \text{Area(R)} + n$, $|\mathcal{Z}_0| + |\mathcal{Z}_1| \leq N^2$.

We prove the following:

Lemma 2 *It is possible to map instances of red-SLP to instances of discrete cover in such a way that red-SLP has a solution if and only if it has a discrete cover. Moreover, the size of an optimal discrete cover is at most four times that of an optimal cover.*

PROOF. Consider an optimal cover for the rectangular region R. We will show how to replace an arbitrary unit square S in a discrete cover with at most four $\mathcal{Z}$-squares S_1, S_2, S_3, S_4 that cover it. Divide S into four closed quadrants Q_1, Q_2, Q_3, Q_4 by drawing the two lines that are parallel to the Cartesian axes and pass from the center of S. Consider one of these quadrants, say Q_i. If a perimeter of a polygon intersects Q_i, then there is a point in $\mathcal{Z}_1$ that lies in Q_i. We take as S_i the unit square centered at such a point of $\mathcal{Z}_1$. If Q_i (its boundary included) does not intersect any polygonal perimeter then it cannot contain a point interior to any of the buildings, because the center of S is not in the interior of any polygon. In this case, there is a point in $\mathcal{Z}_0$ that lies in Q_i. We take as S_i the unit square centered at such a point of $\mathcal{Z}_0$. ∎

In general, we consider the following problem:

Problem 3 (Discrete Station Layout Problem — dis-SLP) *Consider a rectangle R and let $\mathcal{Z}$ be an arbitrary finite set of points in R. Either find a minimum cardinality subset of $\mathcal{Z}$ such that the unit squares centered at its points cover R, or report that no such subset of $\mathcal{Z}$ exists.*

Lemma 2 reduces red-SLP to dis-SLP. Conversely, it is easy to see that dis-SLP can be reduced to red-SLP. To see this consider an instance of dis-SLP defined by a rectangle R and a finite subset $\mathcal{Z}$ of it. Let $\mathcal{P}$ contain a single polygon, "the one obtained from R by deleting the points in $\mathcal{Z}$": by this we mean, to place an arbitrarily small square around each point that is to be deleted and remove the whole square region. This gives a polygon with holes. Then obviously R and $\mathcal{P}$ define an instance of red-SLP which is equivalent to the given instance of dis-SLP.

In the rest of this Section, we examine the dis-SLP instead of red-SLP. As input size N of dis-SLP we take $|\mathcal{Z}| + \text{Area}(R)$. Notice however that when reducing an instance of red-SLP to dis-SLP, the input size of the later is at most the input size of the former squared. This fact, together with the approximation factor mentioned in Lemma 2, makes it necessary to square the time complexity and quadruple the approximation factor of any algorithm for dis-SLP when transformed to an algorithm for red-SLP.

3.3. *Logarithmic approximation algorithm*

In the sequel we give a simple, polynomial time, logarithmic approximation algorithm for dis-SLP by reducing it to the well-known problem SET-COVER (see Garey *al* [4]).

Consider an instance of dis-SLP determined by the rectangle R and the collection of points $\mathcal{Z}$. For each point $p \in \mathcal{Z}$ consider the unit square centered at p. The collection of these unit squares forms a planar subdivision (see Preparata *al* [?] for definition of planar subdivision) of the rectangular region R. Consider the bipartite graph $(\mathcal{A}, \mathcal{Z})$ such that $\mathcal{A}$ is the set of planar rectangular subdomains thus formed. Moreover, for $A \in \mathcal{A}$ and $p \in \mathcal{Z}$, $\{A, p\}$ is an edge if and only if the subdomain A lies entirely inside the unit square centered at p. Now observe that any solution of SET-COVER for the graph $(\mathcal{A}, \mathcal{Z})$ corresponds to a solution of dis-SLP and vice versa. In view of the fact that there are logarithmic approximation algorithms for SET-COVER (e.g., the greedy algorithm in Cormen et al [3]) and the fact that the size of $\mathcal{A}$ is polynomial in the size of $\mathcal{Z}$, we have given a simple polynomial time, logarithmic approximation algorithm to dis-SLP. We improve below this simple result by giving a polynomial, *constant* approximation, dynamic programming algorithm.

3.4. *Constant approximation algorithm*

In this subsection we provide a polynomial time constant approximation algorithm to solve dis-SLP. Our solution is via a reduction to the following problem.

Problem 4 (Discrete Cover Problem — DCP)

Input: *A rectangle R with both height and width of length ≤ 1, and a collection $\mathcal{Z} = \{p_1, p_2, \ldots, p_n\}$ of n points not necessarily all inside the rectangle.*

Output: *Either a minimum cardinality set of unit squares all centered at points of $\mathcal{Z}$ and such that their union covers the rectangle R, or a report that no such set exists.*

We prove the following theorem for DCP.

Theorem 3 *There is an $O(n^8)$ dynamic programming algorithm for DCP, where n is the cardinality of $\mathcal{Z}$.*

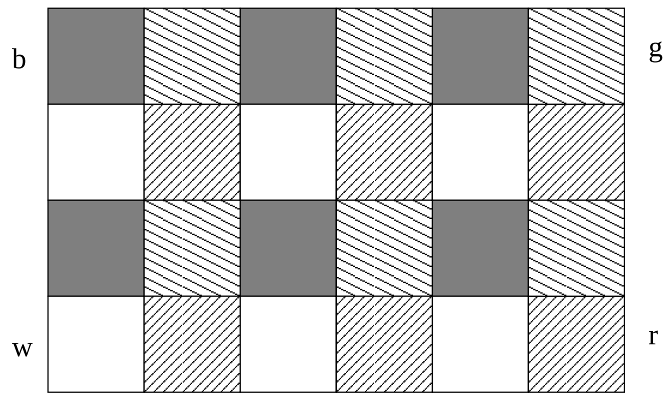

Figure 3: Partitioning the region R like a "checkerboard" of black, white, green and red unit squares.

3.5. *Proof of the main theorem*

Before proving Theorem 3 we indicate how it can be used to find a solution to dis-SLP and hence to SLP. We give the following theorem that is proved by reducing dis-SLP to a number of instances of DCP.

Theorem 4 *There is a polynomial time, constant approximation algorithm for dis-SLP. The constant is at most four.*

PROOF. Partition the rectangular region R as follows. First, left-to-right into vertical strips of width 1, and second top-to-bottom into horizontal strips of width 1 (the last strip in each case may be of length ≤ 1). Now color the resulting rectangles as follows: if the column is odd numbered color the rectangles alternately black and white; if the column is even numbered then color the rectangles alternately green and red, top-to-bottom.

Apply Theorem 3 consecutively to all black, white, green, and red squares. For each subrectangle we get a collection of stations that covers it.

Let b, w, g, r be the total number of stations used to cover black, white, green, and red rectangles, respectively. Let s be the number of stations in an optimal solution. It is clear that $b \leq s$ and $w \leq s$ and $g \leq s$ and $r \leq s$. Since the stations covering different rectangles of the same color do not overlap, it follows that $b+w+g+r \leq 4s$. This completes the proof of Theorem 4. ■

Now from Lemma 2, Theorem 4 and from the discussion at the last paragraph of subsection , we get:

Corollary 1 *There is a polynomial time, constant approximation algorithm for red-SLP and therefore for SLP as well. The approximation factor is 16.*

3.6. *Proof of Theorem 3*

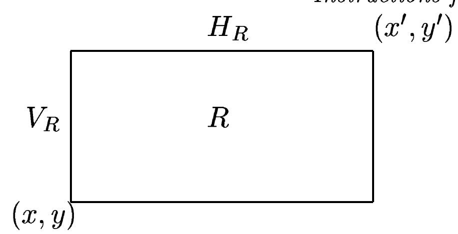

Figure 4: Rectangle $R[x, x', y, y']$ with sides V_R and H_R: the height is $V_R = y' - y$ and the width is $H_R = x' - x$.

The algorithm we give is recursive and is based on dynamic programming. The idea is as follows. We consider the "stations" centered at the given points. For a given rectangle R we consider all possible coverings of this rectangle by stations. We classify the square stations according to how they intersect R, e.g., a square station may either cover R completely, or the left-, right-, down-, up-side of R, or the left-down-, left-up-side, etc. It follows that the number of stations in an optimal solution is determined from solutions to other subrectangles. By scanning the solutions we select the optimal solution to the rectangle R, as will be explained in more detail later on.

3.7. *Classification of the min size cover*

Let $R[x, x', y, y']$ be the axis-parallel rectangle depicted in Figure 4 with lower left corner (x, y) and upper right corner (x', y'). Let the given points be $\mathcal{Z} = \{p_1, p_2, \ldots, p_n\}$ and let S_i denote the unit square centered at p_i. Let $\mathcal{S}$ be the set of all unit squares centered at points of $\mathcal{Z}$. We want to find the minimum size subset $\mathcal{P} \subseteq \mathcal{Z}$ such that the collection of squares centered on elements of $\mathcal{P}$ covers the rectangle R.

We now define the R-Classification of stations. Given a rectangle $R := R[x, x', y, y']$, we classify the squares in $\mathcal{S}$ as follows using the notation C, L,R,U,D for Contains, Left, Right, Up, and Down, respectively. Thus, if x_i and y_i are the coordinates of a point, then $y_i^U, y_i^D, x_i^L, x_i^R$ are the coordinates of the lower left and upper right endpoints of the square centered at the point. We also use the notation LD for the set $L \cap D$, i.e.,

$$\begin{aligned}
C \quad &= \{S_i : S_i \text{ contains } R\} \\
L \quad &= \{S_i : y_i^U \geq y', y_i^D \leq y, x < x_i^R < x', x_i^L \leq x\} \\
LD \quad &= \{S_i : y < y_i^U < y', y_i^D < y, x_i^L < x, x < x_i^R < x'\}
\end{aligned}$$

The other classes R, U, D and LU, RD, RU are defined similarly. The sets L and LD are depicted in Figure 5. Note that these sets are disjoint and their union is equal to $\mathcal{S}$.

Given R and $\mathcal{S}$ let $C^*(\mathcal{S}, R)$ denote the minimum size cover of R by stations from $\mathcal{S}$. We consider the following cases. Each case assumes that the previous case does not hold. With this in mind it is clear that the classification is complete, in

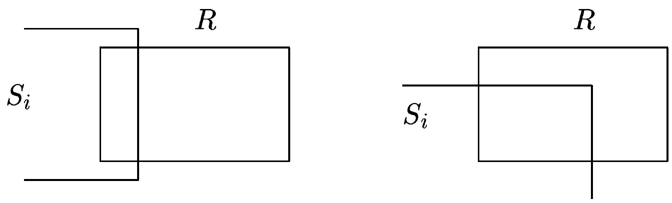

Figure 5: R-Classification: In the left picture $S_i \in L$ and in the right picture $S_i \in LD$.

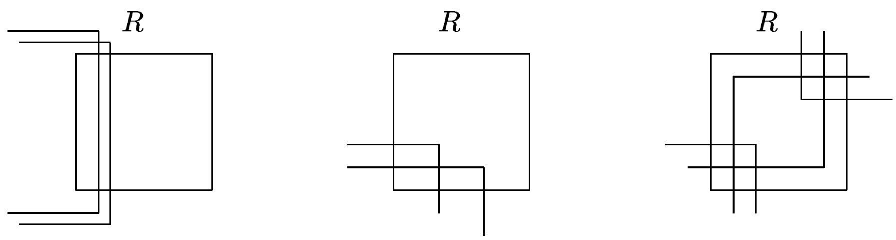

Figure 6: Leftmost figure depicts Case 2L; middle figure depicts Case 3LD, and rightmost figure depicts Case 4.

the sense that $C^*(\mathcal{S}, R)$ must belong to one of the cases below.
Case 1. $C \neq \emptyset$.
In this case $|C^*(\mathcal{S}, R)| = 1$.
Case 2L. $L \neq \emptyset$.
In this case $C^*(\mathcal{S}, R)$ contains exactly one $S_i \in L$ (namely the one farthest to the right which dominates all the other squares in L) (See Figure 6.)
Cases 2R, 2U, 2D. Similar. Next we consider the classes LD, LU, RD, RU. We study only Case 3LD. The other three cases are similar.
Case 3LD. $C^*(\mathcal{S}, R)$ contains at least two squares from LD.
Let the squares of $C^*(\mathcal{S}, R) \cap LD$ be $S_{i_1}, S_{i_2}, \ldots, S_{i_k}$ ordered by ascending x-coordinate. Without loss of generality we may assume that they are also ordered by descending y-coordinate. Indeed, otherwise one of them, say S_{i_j} is dominated by the following square $S_{i_{j+1}}$ in terms of its contribution to covering R, and hence it can be discarded, as shown in Figure 7. Let S_{i_1} and S_{i_2} be any two squares of LD in $C^*(\mathcal{S}, R)$, and let ρ be the upper right intersection point between S_{i_1} and S_{i_2}, $\rho = (\hat{x}, \hat{y}) = (x^R_{i_1}, y^U_{i_2})$ (see Figure 7).
Claim. If case 3LD holds as above, then $C^*(\mathcal{S}, R)$ must contain a square $S_k \in RU$ such that $\rho \in S_k$.
Proof. of the Claim. Consider the infinite set $\{\rho_\epsilon = (\hat{x} + \epsilon, \hat{y} + \epsilon) : \epsilon > 0\}$. There must exist some S_k in $C^*(\mathcal{S}, R)$ that covers the set $\{\rho_\epsilon : 0 < \epsilon < \epsilon_0\}$, for some ϵ_0. We will prove that necessarily $S_k \in RU$ and $\rho \in S_k$. The argument is as follows.

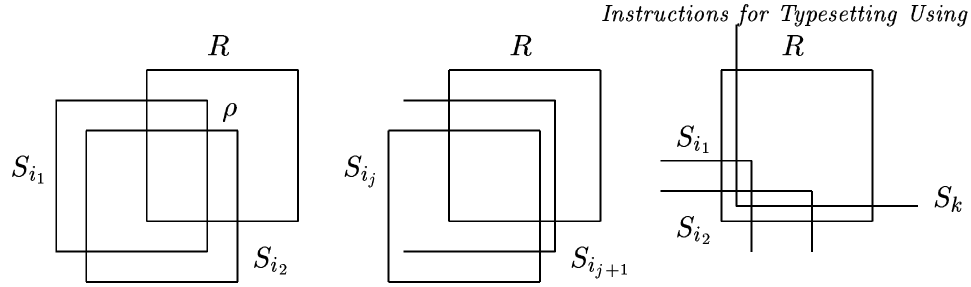

Figure 7: Clasification of $C^*(\mathcal{S}, R)$.

Suppose first that $S_k \in LD$. Then since it does not dominate S_{i_1} or S_{i_2}, it should have appeared between them in the sorted list of $C^*(\mathcal{S}, R) \cap LD$. But this is a contradiction.

Suppose now that $S_k \in RD$ (see Figure 8). As S_k covers all points ρ_ϵ, for $\epsilon < \epsilon_0$, $x_k^L \leq \hat{x}$ and $y_k^U \geq \hat{y}$, hence S_{i_2} is dominated (in terms of covering R) by $S_{i_1} \cup S_k$, contradicting the minimality of $C^*(\mathcal{S}, R)$. If now we suppose $S_k \in LU$ then the proof is similar, with the roles of S_{i_1} and S_{i_2} reversed. Hence the only option left is $S_k \in RU$. Again, as S_k covers ρ_ϵ for every $\epsilon < \epsilon_0$, S_k has coordinates $x_k^L \leq \hat{x}, y_k^D \leq \hat{y}$. Hence $\rho \in S_k$. This completes the proof of the claim.

We note that the same observation as for Case 3LD, holds also for Cases 3LU, 3RD, 3RU. The last case left is the following.

Case 4. $C^*(\mathcal{S}, R)$ contains exactly one square from each of the sets LU, LD, RU, RD.

3.8. *Dynamic programming algorithm*

We are now in a position to use the above R-Classification of squares in order to provide a dynamic programming algorithm computing the minimal number of squares in a covering. An optimal solution is constructed by recursion. The purpose of the previous classification is to establish the fact that all possible cases for the structure of C^* were examined by the algorithm, and no possibility was omitted. Define the sets

$$\begin{aligned} X &= \{x_0^L \leq x_i^L, x_i^R < x_0^R : 1 \leq i \leq n\} \cup \{x_0^L, x_0^R\} \ , \\ Y &= \{y_0^D \leq y_i^D, y_i^U \leq y_0^U : 1 \leq i \leq n\} \cup \{y_0^D, y_0^U\}, \end{aligned}$$

where $x_0^L, x_0^R, y_0^D, y_0^U$ are the coordinates of the original rectangle. For any $x, x' \in X$ and $y, y' \in Y$, let $T(x, x', y, y')$ be the size of the minimum cover of the rectangle $R[x, x', y, y']$ by squares in $\mathcal{S}$. The procedure is the following.

Procedure:

Calculate $T(x, x', y, y')$ for every $x, x' \in X$ and $y, y' \in Y$ by first order, i.e., calculating $T(x, x', y, y')$ only after finishing all rectangles $T(a, a', b, b')$ with both $|a - a'| \leq |x - x'|$ and $|b - b'| \leq |y - y'|$ and $(a, a', b, b') \neq (x, x', y, y')$. In order to calculate $T(x, x', y, y')$ for $R = R[x, x', y, y']$ and $\mathcal{S}$, check systematically through all possibilities for $C^*(\mathcal{S}, R)$.

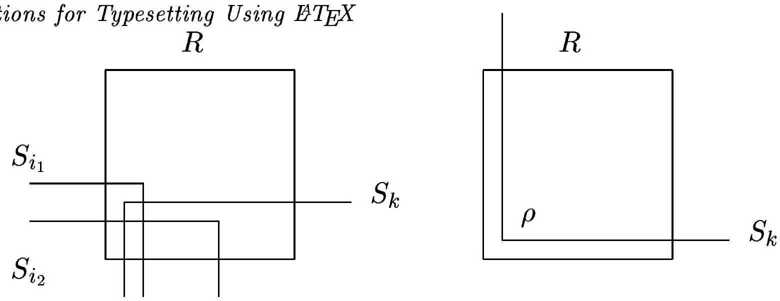

Figure 8: R-Classification of $C^*(\mathcal{S}, R)$.

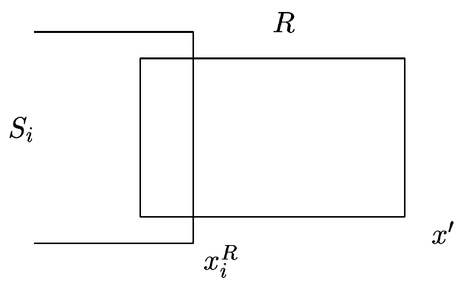

Figure 9: The most dominant square $S_i \in L$.

Case 1. If we are in Case 1, then there should be some S_i that contains R. This is checkable in time $O(n)$.

Case 2L. In this case suppose $L \neq \emptyset$. Go through each $S_i \in L$. For each of these, consult the table concerning the value $t_i = T(x_i^R, x', y, y')$, which is the minimum coverage for $R[x_i^R, x', y, y']$. If such an S_i exists then return $t_i + 1$. Of course it suffices to take the "most dominant" $S_i \in L$, i.e., the one with greatest x_i^R (see Figure 9).*

Cases 2R, 2U, 2D are similar, while Case 4 is easy.

Case 3LD. From the observation we know that in this case we have S_k as in the rightmost picture depicted in Figure 7. Cycle through all choices of $S_{i_1}, S_{i_2} \in LD$

*Now assume that the optimum cover for $R \setminus S_i$ uses some combination of squares that covers also the entire R, say the optimal solution for $R \setminus S_i$ has two elements of the type LU and LD resp., that cover all of the square S_i, and therefore cover the entire R. In this case there is a cover, say $C'(\mathcal{S})$, that uses t_i squares altogether. However, our algorithm will not neglect this correct possibility and will examine it explicitly as part of Cases 3LD, 3LU and 4. Therefore the correct cover will be discovered in due course. In the end the algorithm will take the best cover among all the combinations that were examined, so the best solution wins.

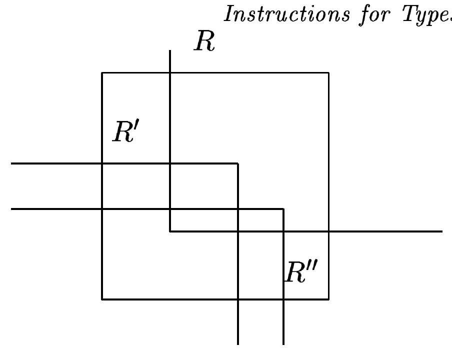

Figure 10: R', R'' are defined by the collection $X \times Y$.

and $S_k \in RU$. If not in the "right shape" ignore. Else (see Figure 10)

$$\begin{aligned}&t' \leftarrow T(R')\\&t'' \leftarrow T(R'')\\&\mathit{Reply}(i_1, i_2, k) \leftarrow t' + t'' + 3\end{aligned}$$

Choose the best of $O(n^3)$ replies $\mathit{Reply}(i_1, i_2, k)$.

Cases 3LU, 3RD, 3RU are similar, while Case 4 is easy.

Combining all these cases we obtain the general procedure for computing $T(x, x', y, y')$ by selecting the best of all replies. The complexity claim follows by noticing that there are $O(n^4)$ subproblems and that Case 4 requires time $O(n^4)$. This completes the proof of Theorem 3. ■

4. The Case of "Thin" Buildings

In this section we provide a linear time, constant approximation algorithm when the buildings satisfy certain width contraints. The approxination constant in this case is 2.

The horizontal h (respectively, vertical v) width of an polygon with edges parallel to one of the Cartesian axes is the maximum length of a horizontal (respectively, vertical) line segment that lies inside the polygon.

Theorem 5 *If either $h + 2r_0 \leq 2r$ or $v + 2r_0 \leq 2r$ then there is a linear time algorithm for finding a solution to SLP, such that the number of stations is at most 2 times the optimal.*

PROOF. We work with red-SLP, i.e., we assume that $r_0 = 0$ and $r = 1/2$. We use a technique of Hochbaum *et al* [6]. Divide the plane into strips of width 1 by drawing lines parallel to the x-axis (the last one may have width ≤ 1).

Consider the ith strip, i.e., the one delimited by the lines L_i, L_{i+1} (in order to have pairwise disjoint strips, assume that the ith strip contains L_i but not L_{i+1}). We can layout an optimal number of stations to cover this strip in a fashion that avoids the interiors of all buildings by proceeding in analogy to the proof of Theorem 1, i.e.,

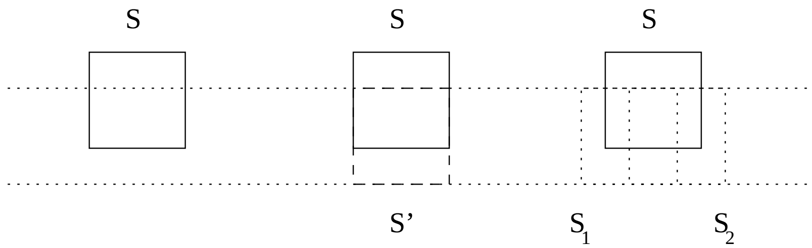

Figure 11: Replacing the original station first with one and then with two stations.

moving along a middle line of the strip and suitably placing unit squares centered on it. Let s_i be the number of stations used in the ith strip. We show that the total number of stations, $\sum_i s_i$, thus used does not exceed two times the optimal. Indeed, let $\mathcal{S}$ be a collection of stations of minimum size covering the buildings. For each i consider the set $\mathcal{S}_i$ of unit squares whose centers lie within the ith strip. It is clear that the sets $\mathcal{S}_i$ are disjoint. Moreover, $\mathcal{S} = \bigcup_i \mathcal{S}_i$. It follows that $|\mathcal{S}| = \sum_i |\mathcal{S}_i|$.

Next we show that $s_i \leq 2|\mathcal{S}_i|$. Indeed, first notice that the points in the middle line of the ith strip can only be covered by squares in $\mathcal{S}_i$. Now slide vertically each square S in $\mathcal{S}_i$, keeping its x-coordinate fixed, until it fits exactly inside the ith strip (recall that the strips have width 1, i.e., equal to the edge-length of the squares). Let S' be the slid square. Since $\mathcal{S}$ covers the whole rectangular area (points in the middle lines of strips included), by the previous remark the collection of squares S' obtained by sliding the elements of $\mathcal{S}_i$ cover the ith strip .

However, in general a slid square S' may not lie in "legal" position because its center p' may lie in the interior of a building. Thus we slide the point p' (keeping the same y-coordinate) to the left and to the right until we determine two "legal" positions p_1, p_2, respectively, and squares S_1 and S_2 centered on p_1, p_2, respectively, such that $S' \subseteq S_1 \cup S_2$. Both S_1, S_2 lie in the strip delimited by the lies L_i and L_{i+1}. This is possible because of our hypothesis $h \leq 1$. This is depicted in Figure 11.

Since the arrangement on each strip is optimal, it follows that $s_i \leq 2|\mathcal{S}_i|$. This completes the proof of Theorem 5. ∎

5. Conclusion and Open Problems

We have considered the problem of covering a rectangular region containing axis-parallel buildings with stations in the presence of location constraints. We have given constant as well as logarithmic approximation polynomial time algorithms for solving the problem.

It is an open problem to determine whether or not finding an optimal solution can be done in polynomial time. We do not know if the problem is either NP hard or difficult to approximate. Several interesting open problems remain. One is to assume there is an upper bound on the number of stations permitted to cover a given point in the region. (As we indicated in subsection , such coverage may not always exist.) Another is to consider stations with varying amounts of available power, as well as studying the impact of having the option to transmit with lower power in order to keep the signal strength at a building sufficiently low.

Acknowledgements

Many thanks to the anonymous referee for detailed and useful comments that improved the presentation of the paper.

Research of Prosenjit Bose and Evangelos Kranakis supported in part by NSERC (Natural Sciences and Engineering Research Council of Canada) and MITACS project CANCCOM (Complex Adaptive Networks for Computing and Communication) grants.

A preliminary version of this paper has appeared in the proceedings of 10th International Symposium on Algorithms and Computation (ISAAC '99), December 16-18, Chennai, India, A. Aggarwal and C. Pandu Rangan, eds., Springer Lecture Notes in Computer Science, pp. 269-278, Vol. 1741, 1999.

References

1. L. I. Aupperle, H. E. Conn, J. M. Kell, and J. O'Rourke, "Covering Orthogonal Polygons with Squares", in: Proceedings of the 26th Annual Allerton Conference on Communication, Control, and Computing, 97–106, 1988.
2. R. Bar-Yehuda, and E. Ben-Chanoch, "An $O(N \log^* N)$ Time Algorithm for Covering Simple Polygons with Squares", in: Proceedings of the 2nd Canadian Conference on Computational Geometry, 186–190, 1990.
3. T. H. Cormen, C. E. Leiserson, and R. L. Rivest, "An Introduction to Algorithms", MIT Press, 1990.
4. M. R. Garey and D. S. Johnson, Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman and Company, 1979.
5. N. A. DePano, Y. Ko, and J. O'Rourke, "Finding Largest Equilateral Triangles and Squares", in: Proceedings of the 25th Annual Allerton Conference on Communication, Control, and Computing, 869–878, 1987.
6. D. S. Hochbaum and W. Maass, "Approximation Algorithms for Covering and Packing Problems in Image Processing and VLSI", J. ACM 32, 130–136, 1985.
7. J.M. Horowitz, "Radio Freakquency", Time Magazine 157(19), May 14 2001.
8. J. C. Lin, "Biological Aspects of Mobile Communication Data", Wireless Networks 3(6), 439–453, 1997.
9. F. P. Preparata and M. I. Shamos, Computational Geometry: An Introduction, Springer-Verlag, Texts and Monographs in Computer Science, 1985.
10. http://www.metricom.com/ricochet_advantage