

Efficient Approximation Schemes for Geometric Problems?*

Dániel Marx

Department of Computer Science and Information Theory,
Budapest University of Technology and Economics,
Budapest, H-1521, Hungary
dmarx@cs.bme.hu

Abstract. An EPTAS (efficient PTAS) is an approximation scheme where ϵ does not appear in the exponent of n , i.e., the running time is $f(\epsilon) \cdot n^c$. We use parameterized complexity to investigate the possibility of improving the known approximation schemes for certain geometric problems to EPTAS. Answering an open question of Alber and Fiala [2], we show that MAXIMUM INDEPENDENT SET is W[1]-complete for the intersection graphs of unit disks and axis-parallel unit squares in the plane. A standard consequence of this result is that the $n^{O(1/\epsilon)}$ time PTAS of Hunt et al. [11] for MAXIMUM INDEPENDENT SET on unit disk graphs cannot be improved to an EPTAS. Similar results are obtained for the problem of covering points with squares.

1 Introduction

We say that an optimization problem admits a polynomial-time approximation scheme (PTAS) if for every $\epsilon > 0$ there is a polynomial-time algorithm with relative error at most ϵ . A PTAS is an *efficient polynomial-time approximation scheme* (EPTAS) if this family of approximation algorithms is uniformly polynomial: for every $\epsilon > 0$, the running-time is $f(\epsilon) \cdot n^c$, where f is an arbitrary function of ϵ , and c is a constant independent of ϵ . For example, Arora [3] presented an $n^{O(1/\epsilon)}$ time PTAS for Euclidean TSP, which is not an EPTAS. However, in the journal version of the paper [4], the running-time of the algorithm is improved to $n \cdot \log^{O(1/\epsilon)} n = 2^{O(1/\epsilon^2)} \cdot n^2$, hence the problem admits an EPTAS.

Whenever a problem admits a PTAS, it should be examined if the algorithm can be improved to an EPTAS. The motivation comes from the observation that a polynomial-time algorithm is not really practical if the degree is larger than 3. Therefore, a $O(n^{2/\epsilon})$ time PTAS is not practical even for 20% error. In fact, the situation is much worse than that: as pointed out in [7], most approximation schemes in the literature have very high degrees even for 20% error. For example, the running time of the PTAS of [9] for finding the maximum weighted independent set in the intersection graph of disks is $O(n^{523804})$ for 20% error.

* Research is supported in part by grants OTKA 44733, 42559 and 42706 of the Hungarian National Science Fund.

Parameterized complexity gives us useful tools to investigate the question whether a PTAS can be improved to an EPTAS. Parameterized complexity deals with problems where the input instances have a distinguished part k called the parameter. For example, in the MAXIMUM CLIQUE problem the input is of the form (G, k) , where k is the size of the clique to be found. Usually we are considering problems that are polynomial-time solvable for every fixed value of k : for example, MAXIMUM CLIQUE can be solved by checking all the $O(n^k)$ size k sets. However, in parameterized complexity we are interested in the question whether there is a *uniformly polynomial-time* algorithm for the problem, that is, whether it can be solved in $f(k)n^c$ time, where f depends only on k and c is a constant independent of k . It turns out that several NP-hard problems, such as MINIMUM VERTEX COVER, LONGEST PATH, and TRIANGLE PACKING can be solved in uniformly polynomial time. Such problems are called *fixed-parameter tractable* (FPT). On the other hand, for MAXIMUM CLIQUE, MINIMUM DOMINATING SET, and several other problems, no uniformly polynomial algorithm is known. W[1]-hardness is the parameterized complexity analog of NP-hardness: by showing that a problem is W[1]-hard, we prove that the problem is not fixed-parameter tractable (under standard complexity-theoretic assumptions). For more background, the reader referred to the monograph of Downey and Fellows [8].

In [5] and [6], it is noted that an approximation algorithm with relative error $1/(2k)$ can decide whether the optimum is k . Therefore, an EPTAS with running time $f(\epsilon)n^c$ immediately implies that the parameterized version of the problem is fixed-parameter tractable: the EPTAS gives an $f(1/(2k))n^c$ time algorithm for the problem. This means that by proving that a problem is W[1]-hard, we can show that the problem is unlikely to have an EPTAS.

There are many optimization problems involving geometric objects in the plane that admit a PTAS. For example, [11] presents a PTAS for finding the maximum number of pairwise independent disks in a collection of unit disks. The reason for the abundance of approximation schemes for geometric problems is that in many cases, shifting and layering techniques can be used to reduce the problem to small subproblems that can be solved by brute force. In this paper we investigate whether it is possible to give an EPTAS for such problems.

For a set V of geometric objects, the *intersection graph* of V is a graph with vertex set V where two vertices are connected if and only if the two objects have non-empty intersection. Intersection graphs of disks, rectangles, line segments and other objects arise in applications such as facility location [14], frequency assignment [12], and map labeling [1]. Maximum independent set is NP-hard for the intersection graphs of unit disks, but admits an $n^{O(1/\epsilon)}$ time PTAS [11]. Alber and Fiala [2] considered the special case of λ -precision unit disk graphs, where the distance between the centers of disks have to be at least λ . For fixed λ , they gave a $O(2^{\sqrt{k}} + n)$ algorithm for finding an independent set of size k , which shows that the problem is FPT in this special case. The parameterized complexity of the problem without the precision restriction remained an open question. Here we answer this question by showing that maximum independent set is W[1]-complete for unit disk graphs. This result has two implications. First,

it shows that (unless $\text{FPT} = \text{W}[1]$) the uniformly polynomial algorithm of [2] for λ -precision unit disk graphs cannot be extended to general unit disk graphs. Furthermore, the PTAS of [11] cannot be improved to an EPTAS; the expression $1/\epsilon$ cannot be taken out of the exponent of n in the running time. Clearly, the same conclusion holds for the $n^{O(1/\epsilon^4)}$ PTAS of [9] that solves the more general problem of weighted independent set for disks with arbitrary diameter.

On the positive side, if we have a PTAS for a problem whose parameterized version is FPT, then this might indicate that the PTAS can be improved to an EPTAS. We present two examples for this situation. As shown in [2], the MAXIMUM INDEPENDENT SET problem is FPT for λ -precision unit disk graphs, and by [11], there is a linear-time EPTAS for such graphs. Furthermore, unlike MAXIMUM INDEPENDENT SET, the MINIMUM VERTEX COVER problem is FPT for every graph, hence it might be possible that there is an EPTAS for this problem on unit disk graphs. This is indeed so: in Section 3, we show that the PTAS of [11] can be improved to a linear-time EPTAS.

In Section 4, we study the problem of covering a given set of points by as few squares as possible. This problem is motivated by applications in image processing. Hochbaum and Maass [10] presented an $n^{O(1/\epsilon^2)}$ time PTAS for the problem. By proving that the corresponding parameterized problem is $\text{W}[1]$ -hard, we show that this PTAS cannot be improved to an EPTAS.

2 Maximum Independent Set

We prove that MAXIMUM INDEPENDENT SET remains $\text{W}[1]$ -complete when restricted to the intersection graphs of unit disks and unit squares. First we prove $\text{W}[1]$ -completeness for unit squares by a parameterized reduction from the MAXIMUM CLIQUE problem. A parameterized reduction transforms an instance (I, k) to an instance (I', k') where k' depends only on k , but not on I (ordinary reductions used to show NP-completeness usually do not have this property). As shown in Theorem 3, essentially the same reduction can be used for unit disks.

Theorem 1. *MAXIMUM INDEPENDENT SET is $\text{W}[1]$ -hard for the intersection graphs of axis-parallel unit squares in the plane.*

Proof. The reduction is from parameterized MAXIMUM CLIQUE. We have to determine whether the given graph G contains a clique of size k . For convenience, we assume that the number of vertices and the number of edges are both n .¹ The squares are open, two squares that share only a boundary do not intersect.

Set $\epsilon := 1/n^2$. The squares constructed in the reduction are partitioned into k' blocks, where k' depends only on k . Each block has a position (x, y) , which is a pair of integers. If a square belongs to the block at (x, y) , then the coordinates of its lower left corner are of the form $(x + i\epsilon, y + j\epsilon)$ for some integers $1 \leq i, j \leq n$. Therefore, a block can contain at most n^2 squares; a block containing all n^2 of them will be called a *complete block*. If the horizontal coordinate of a square is

¹ This can be achieved by adding/deleting isolated vertices and acyclic components.




Fig. 1. The structure of the core

of the form $x + i\epsilon$, then we say that the *horizontal offset* of the square is i . The vertical offset is similarly defined. At most one square can be selected from each block, hence every independent set in G' has size at most k' . Furthermore, every size k' independent set contains exactly one square from each block.

The constructed instance consists of two parts: the core and the wrap-around machinery. The core is illustrated in Figure 1. There are k rows with $3\binom{k}{2}$ blocks in each, and there are $\binom{k}{2}$ columns with $3k$ blocks in each. The columns are indexed by the two element subsets of $\{1, 2, \dots, k\}$, i.e., each index is a pair (j_1, j_2) with $1 \leq j_1 < j_2 \leq k$. The block in the intersection of row i and column (j_1, j_2) is denoted by $B_{i,(j_1,j_2)}$. The set of squares contained in block $B_{i,(j_1,j_2)}$ is defined as follows. If $i \neq j_1$ and $i \neq j_2$, then the block $B_{i,(j_1,j_2)}$ is a complete block. Let e_1, e_2, \dots, e_n be the edges of G , and let $v^{(1)}(e_j) < v^{(2)}(e_j)$ be the two vertices of edge e_j . The block $B_{j_1,(j_1,j_2)}$ contains n squares: for each $1 \leq j \leq n$, it contains a square with horizontal offset $v^{(1)}(e_j)$ and vertical offset j . Similarly, $B_{j_2,(j_1,j_2)}$ contains a square with horizontal offset $v^{(2)}(e_j)$ and vertical offset j .

There are also blocks that are in one of the rows and in none of the columns, and there are blocks that are in one of the columns and in none of the rows. All these blocks (shown in lighter color in Fig. 1) are complete blocks.

Assume that one square is selected from each block of the core in such a way that they are pairwise non-intersecting. Such a solution will be called a *standard solution* if every square selected from row i has the same horizontal offset x_i and every square from column (j_1, j_2) has the same vertical offset $y_{(j_1,j_2)}$.

Lemma 2. *The core has a standard solution if and only if G has a size k clique.*

Proof. Assume that the core has a standard solution with horizontal offsets x_i and vertical offsets $y_{(j_1,j_2)}$. We claim that the numbers x_i correspond to a clique of size k in G . Suppose that, on the contrary, for some $1 \leq j_1 < j_2 \leq k$ the pair $x_{j_1}x_{j_2}$ is not an edge of G (including the possibility that $x_{j_1} = x_{j_2}$). Consider the square selected from block $B_{j_1,(j_1,j_2)}$, it has horizontal offset x_{j_1} and vertical offset $y' = y_{(j_1,j_2)}$. By construction, this means that $v^{(1)}(e_{y'}) = x_{j_1}$, i.e., the first vertex of edge $e_{y'}$ is x_{j_1} . Similarly, by considering the square selected from $B_{j_2,(j_1,j_2)}$, the second vertex of $e_{y'}$ is x_{j_2} , hence $x_{j_1}x_{j_2}$ is an edge.

Now assume that x_1, \dots, x_k is a clique of size k in G , with edge $e_{(j_1,j_2)}$ being the edge connecting x_{j_1} and x_{j_2} . Let the horizontal offset in row i be x_i , and let the vertical offset in column (j_1, j_2) be $e_{(j_1,j_2)}$. The construction ensures that we can select such a square from block $B_{i,(j_1,j_2)}$. Consider the four complete blocks next to $B_{i,(j_1,j_2)}$, and select a square from each of them that has the same offsets

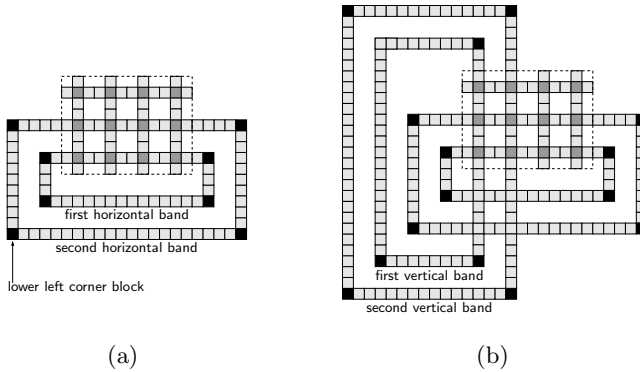


Fig. 2. (a) The first two horizontal wrap-around bands. (b) The first two horizontal and vertical wrap-around bands.

as the one selected from $B_{i,(j_1,j_2)}$. The square from $B_{i,(j_1,j_2)}$ and the four squares around it are pairwise non-intersecting, thus we get a (standard) solution. \square

The wrap-around machinery consists of k horizontal bands and $\binom{k}{2}$ vertical bands. The i -th horizontal band ensures that the selected squares in the i -th row have the same horizontal offset. Each horizontal band connects the first and the last block of the row, as shown in Figure 2a. Notice that the distance is at least two between two bands. Apart from the four corner blocks, each block is a complete block. The upper left and the lower right corner blocks have the same structure: for every $1 \leq i \leq n$, they contain the squares with horizontal offset i and vertical offset i . The upper right and the lower left corner blocks contain the squares with horizontal offset i and vertical offset $n - i + 1$ ($1 \leq i \leq n$).

We show that if one square is selected from each block such that they are independent, then the squares in row i have the same horizontal offset. Let ℓ_i (resp., r_i) be the horizontal offset of the first (resp., last) square of row i . The horizontal offset of the second square in row i is at least ℓ_i , otherwise it would conflict with the first square. Continuing this argument, we get that $\ell_i \leq r_i$, and the horizontal offset of every square in row i is between ℓ_i and r_i . Therefore, it is sufficient to show $r_i \leq \ell_i$. Consider the horizontal segment of the band extending from the last block of row i . By the same argument as above, the horizontal offset of each selected square is at least r_i . In particular, this is true for the upper right corner block. By construction, this means that the vertical offset of the square in that corner block is at most $n - r_i + 1$. This implies that in the right vertical segment of the band every block has a vertical offset of at most $n - r_i + 1$. Therefore, the horizontal offset of the square in the lower right corner block is at most $n - r_i + 1$. Continuing further in a similar way, it follows that the square in the lower left corner has vertical offset at least r_i , the square in the upper left corner has horizontal offset at least r_i , hence the horizontal offset of the first block of row i is also at least r_i . On the other hand, it is easy to see that if the square in the first and last block of row i have the same horizontal offset, then we can select one square from each block of the i -th horizontal band.

The $\binom{k}{2}$ vertical bands are defined analogously, band (j_1, j_2) ensures that the vertical offset is the same for every square selected from column (j_1, j_2) (see Figure 2b.) Notice that we reuse some blocks of the horizontal bands when a vertical band crosses a horizontal band, but that does not modify our conclusion that the bands enforce a standard solution for the core. The only thing that has to be carefully examined is whether a standard solution of the core can be extended to the horizontal and vertical bands simultaneously. Consider the block B at the intersection of horizontal band i and vertical band (j_1, j_2) . The horizontal (resp., vertical) band determines the horizontal (resp., vertical) offset of the square in this block. There are 4 other blocks next to block B , from each of them we select a square with the same horizontal and vertical offset as the one selected from B . The selected squares do not intersect each other, and they do not intersect the other squares on their band. This can be done independently for every intersection, since their distance is at least two blocks.

The reduction constructs k' blocks, where k' depends only on k . We have shown that there are k' independent squares if and only if there is a standard solution for the core, or equivalently, if the graph G has a clique of size k . \square

The same reduction shows hardness for unit disks. For convenience, we present the proof for unit-diameter disks instead of unit-radius disks.

Theorem 3. MAXIMUM INDEPENDENT SET is $W[1]$ -hard for the intersection graphs of unit disks in the plane.

Proof. The same reduction works as in Theorem 1: a square with lower left corner at $(x + i\epsilon, y + j\epsilon)$ can be replaced by a unit disk with center $(x + i\epsilon, y + j\epsilon)$. In the reduction, only the following two properties of the squares were used:

1. If we select a square with horizontal offset i_1 from the block at position (x, y) , and we select a square with horizontal offset i_2 from the block at position $(x + 1, y)$, then they intersect if and only if $i_1 > i_2$.
2. If we select a square with vertical offset j_1 from the block at position (x, y) , and we select a square with vertical offset j_2 from the block at position $(x, y + 1)$, then they intersect if and only if $j_1 > j_2$.

We show that the same properties hold for (open) unit disks. Consider two disks with centers $(x + i_1\epsilon, y + j_1\epsilon)$ and $(x + 1 + i_2\epsilon, y + j_2\epsilon)$. Clearly, if $i_1 \leq i_2$, then their distance is at least 1. On the other hand, if $i_1 > i_2$ then their distance is

$$\begin{aligned} & \sqrt{(1 + (i_2 - i_1)\epsilon)^2 + (j_2 - j_1)^2\epsilon^2} \\ & \leq \sqrt{(1 - \epsilon)^2 + n^2\epsilon^2} = \sqrt{1 - 2\epsilon + (n^2 + 1)\epsilon^2} < 1, \end{aligned}$$

if $\epsilon \leq 1/n^2$. Property 2 can be shown similarly. \square

The reduction constructs instances where the centers are arbitrarily close to each other. Therefore, the proof does not work for λ -precision unit disk graphs. This is not surprising, since it is shown in [2] that for every fixed $\lambda > 0$, MAXIMUM INDEPENDENT SET is fixed-parameter tractable for λ -precision unit disk

graphs. This opens the possibility that (unlike the general case) the problem restricted to λ -precision unit disk graphs admits an EPTAS. Indeed, [11] shows that there is a linear-time EPTAS for every fixed λ .

3 Minimum Vertex Cover

After giving a PTAS for MAXIMUM INDEPENDENT SET in unit disk graphs, Hunt et al. [11] briefly discuss a similar PTAS for the MINIMUM VERTEX COVER problem. From the parameterized complexity point of view, MAXIMUM INDEPENDENT SET and MINIMUM VERTEX COVER are very different: the first problem is W[1]-hard (even for unit disk graphs, Section 2), while the latter problem is fixed-parameter tractable for *every* graph (the current best algorithm is presented in [13]). Therefore, we cannot prove a MINIMUM VERTEX COVER analog of Theorem 3, which raises the possibility that the PTAS of [11] for MINIMUM VERTEX COVER can be improved to an EPTAS. We show here that some simple ideas are sufficient to turn this PTAS into a linear-time algorithm.

Let D be a set of unit-diameter disks. In the first phase of the algorithm, we ensure that every point of the plane is contained in at most $1/\epsilon$ disks. If point p is contained in more than $1/\epsilon$ disks, then add these disks into the set S , and remove them from D . We repeat this until no such p can be found, let D_0 be the set of remaining disks. We claim that the set S together with a $(1 + \epsilon)$ -approximation of the vertex cover for D_0 gives a $(1 + \epsilon)$ -approximate vertex cover for D . This follows from the fact that whenever we add to S the $\ell \geq 1/\epsilon + 1$ disks containing some point p , then at least $\ell - 1$ of these disks have to appear in every vertex cover. Hence S itself is a $(1 + \epsilon)$ -approximate vertex cover of S .

The linear-time EPTAS of [11] for MAXIMUM INDEPENDENT SET in λ -precision unit disk graphs is based on the observation that a constant-sized rectangle can contain at most a constant number of disks. If we perform the first phase of the algorithm, then this property will hold for our instance. Therefore, we can give a linear-time PTAS for MINIMUM VERTEX COVER that is similar to [11–Theorem 5.2] (we omit the details).

Theorem 4. *There is a $2^{O(1/\epsilon^2)} \cdot n$ time EPTAS for MINIMUM VERTEX COVER in unit disk graphs.* \square

4 Covering Points with Squares

Hochbaum and Maass [10] presented a PTAS for the problem of covering n given points in \mathbb{R}^d with the minimum number of d -dimensional unit-diameter balls or d -dimensional rectilinear blocks. As a special case, their result gives an $n^{O(1/\epsilon^2)}$ time approximation scheme for covering n points in the plane by unit squares. We show that it is unlikely that this PTAS can be improved to an EPTAS, since the parameterized version of the problem is W[1]-hard.

Theorem 5. COVERING POINTS WITH SQUARES *is* W[1]-hard.

Proof. The proof is similar in structure to the proof of Theorem 1. The reduction is from MAXIMUM CLIQUE; the k vertices are selected by k rows, and the $\binom{k}{2}$ edges are selected by $\binom{k}{2}$ columns. Horizontal and vertical bands are used to ensure the consistency of the rows and columns (see below for details). In Theorem 1, the structure of the graph was encoded by the squares available in certain blocks. In COVERING POINTS WITH SQUARES, we cannot prescribe which squares can be used in a solution, hence a more delicate construction is required to ensure that the selected vertices and edges form a correct solution.

We will use the directions east (increasing x coordinate), north (increasing y coordinate), northeast, etc. The directions will be abbreviated as E, N, NE, etc. For convenience, we assume that the squares are closed on west and south, and open on east and on north. The SW- and the SE-corner points belong to the square, but the NW- and the NE-corners do not. (It can be shown that the problem has the same complexity with open, closed, and half-open squares).

Let G be the graph where we have to find a size k clique. For convenience, we assume that the number of edges and vertices are both n in G . Set $\epsilon := 1/n^2$. Every point constructed in the reduction has coordinates that are integer multiples of ϵ . Therefore, it can be assumed that in a solution the coordinates of the corners of each square are integer multiples of ϵ .

We use the points to construct *blocks*, *connectors*, and *testers*. If there is a *block* at (x, y) , then this means that we add 5 points $(x+0.5, y+0.5)$, $(x+n\epsilon, y+0.5)$, $(x+1-n\epsilon-\epsilon, y+0.5)$, $(x+0.5, y+n\epsilon)$, $(x+0.5, y+1-n\epsilon-\epsilon)$. These 5 points are called the central, W, E, S, N *control points* of the block, respectively.

The problem parameter k' in the constructed instance of COVERING POINTS WITH SQUARES is equal to the number blocks. It can be shown (details omitted) that the only way k' squares can cover the control points of k' blocks is that if every square corresponds to some block:

Proposition 6. *For each block, there is a unique square in the solution that covers all five control points of the block. □*

Therefore, the SW-corner of the square of block (x, y) has coordinates $(x+i\epsilon, y+j\epsilon)$ for some integers $-n \leq i \leq n$, $-n \leq j \leq n$. These two integers are called the *horizontal and vertical offsets* of block (x, y) .

We will add *boundary points* to some of the blocks. There are four types of boundary points: N, S, E, W. Whenever we add a N (S etc.) boundary point to a block, then it will be true that there is no N (S etc.) neighbor of the block. Adding boundary points to the block at (x, y) will be done as follows:

- The N-boundary point is at $(x+0.5, y+1)$. It ensures that the vertical offset of the block is positive (recall that the squares are open on north).
- The S-boundary point is at $(x+0.5, y)$. It ensures that the vertical offset of the block is not positive.

The E-boundary (resp., W-boundary) points are defined analogously, they ensure that the horizontal offset is positive (resp., not positive).

A *connector* is a set of points whose job is to ensure that certain relations hold between the offsets of two neighboring blocks. A horizontal connector between

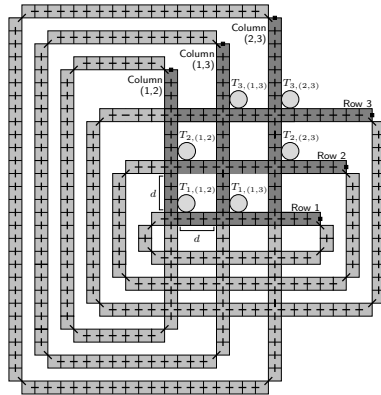


Fig. 3. Structure of the constructed instance in the proof of Theorem 5

blocks (x, y) and $(x + 1, y)$ consists of the $2n$ points $(x + 1 + i\epsilon, y + 0.5)$ ($-n \leq i \leq n - 1$). These points can be covered only by the squares of blocks (x, y) and $(x + 1, y)$. These two squares cover the connector if and only if the horizontal offset of block (x, y) is not smaller than the horizontal offset of block $(x + 1, y)$.

Similarly, the vertical connector between blocks (x, y) and $(x, y + 1)$ consists of the points $(x + 0.5, y + 1 + i\epsilon)$ ($-n \leq i \leq n - 1$). These points ensure that the vertical offset of (x, y) is not the smaller than the vertical offset of $(x, y + 1)$.

A diagonal connector between blocks (x, y) and $(x + 1, y + 1)$ consists of the points $(x + 1 + i\epsilon, y + 1 + i\epsilon)$ ($-n \leq i \leq n - 1$). These points ensure that if block (x, y) has offsets i_1, j_1 , and block $(x + 1, y + 1)$ has offsets i_2, j_2 , then $i_2, j_2 \leq \min(i_1, j_1)$. The blocks $(x, y + 1)$ and $(x + 1, y)$ can be connected in a similar way. In this case, if i_1, j_1 are the offsets of $(x, y + 1)$, and i_2, j_2 are the offsets of $(x + 1, y)$, then $i_2 \leq \min(i_1, -j_1)$ and $j_2 \geq \max(-i_1, j_1)$ follows.

Figure 3 shows the structure of the constructed instance of COVERING POINTS WITH SQUARES. As in Theorem 1, there are k rows (shown by darker blocks) that correspond to the k vertices of the clique, and there are $\binom{k}{2}$ columns (also shown in dark) that correspond to the $\binom{k}{2}$ edges of the clique. The rows are indexed from 1 to k , while the column indexes are pairs (j_1, j_2) ($1 \leq j_1 < j_2 \leq k$). The connector gadgets connecting the neighboring blocks are shown by short line segments in the figure. There are $2\binom{k}{2}$ tester gadgets (shown by circles in Figure 3): for every (j_1, j_2) ($1 \leq j_1 < j_2 \leq k$), there is a tester gadget $T_{j_1, (j_1, j_2)}$ connected to both row j_1 and column (j_1, j_2) , and there is a tester gadget $T_{j_2, (j_1, j_2)}$ connected to both row j_2 and column (j_1, j_2) . The distance between the rows/columns is d blocks, where d is sufficiently large to ensure that there is enough space for the tester gadgets between the rows and columns (e.g., $d = 20$ is enough).

For each row, the leftmost and the rightmost blocks are connected by a horizontal band. As in the proof of Theorem 1, this band ensures that in every solution, the horizontal offset is the same for every block of the row. If we follow what requirements the connectors prescribe on the adjacent blocks, then it turns




Fig. 4. The tester gadget

out that the horizontal offset of the rightmost block cannot be smaller or larger than the horizontal offset of the leftmost block. Therefore, the same horizontal offset has to appear in every block of the row. Similarly, the vertical bands ensure that the vertical offset is the same for every block of a given column.

We add an E-boundary point to the rightmost block of each row (see the small dot in Fig. 3). Similarly, we add a N-boundary point to the topmost block of each column. These points ensure that the horizontal (resp., vertical) offset is between 1 and n for every row (resp., column).

Given a solution to the constructed instance of COVERING POINTS WITH SQUARES, we interpret the horizontal offset of row i as the index of the i -th vertex of the clique, and the vertical offset of column (j_1, j_2) as the index of the edge e_{j_1, j_2} connecting the j_1 -th and j_2 -th vertices. The tester gadgets ensure that this interpretation gives a consistent solution: the two end points of the edge given by column (j_1, j_2) are the vertices given by rows j_1 and j_2 . More precisely, gadget $T_{j_1, (j_1, j_2)}$ ensures that the smaller end of e_{j_1, j_2} is the j_1 -th vertex of the clique, and $T_{j_2, (j_1, j_2)}$ ensures that the larger end of e_{j_1, j_2} is the j_2 -th vertex.

The tester gadget $T_{j_1, (j_1, j_2)}$ connected to row j_1 and column (j_1, j_2) is constructed as follows (the description of $T_{j_2, (j_1, j_2)}$ is similar). The gadget consists of 33 blocks and it is arranged as shown in Fig. 4. Besides these new blocks, we add some additional points to the rows and columns as well. A W-boundary point is added to the W-neighbor of block r and to the W-neighbor of block s (see Fig. 4). A S-boundary point is added to the S-neighbors of blocks t and u .

Let (x, y) be the coordinates of the SE-corner of block t . Let us add the points $(x + \ell\epsilon, y + \epsilon)$ ($1 \leq \ell \leq n$). The S-neighbor of block t has a S-boundary point, thus only block t and its SE-neighbor, z , can cover these new points. This means that the horizontal offset of block t cannot be smaller than the horizontal offset of row j_1 . On the other hand, if the horizontal offset of row j_1 is α , then z can cover $(x + \ell\epsilon, y + \epsilon)$ for every $\ell \geq \alpha$ if we set the vertical offset of z to at least 2. In this case, t has to cover $(x + \ell\epsilon, y + \epsilon)$ only for $\ell < \alpha$, hence if the horizontal offset of t is not smaller than α , then all the points are covered.

In a similar way, we force the vertical offset of block s to be at least as large as the vertical offset of column (j_1, j_2) . For this purpose, if (x, y) are the coordinates of the NW-corner of block s , then we add the points $(x + \epsilon, y + \ell\epsilon)$ ($1 \leq \ell \leq n$). The argument is the same as in the previous paragraph.

There is a diagonal connector between blocks z and u . However, this connector is slightly different from the one defined above. Let (x, y) be the coordinates of the NE-corner of block z . The connector consists of the points $(x + (\ell + 1)\epsilon, y + \ell\epsilon)$ for $-n \leq \ell \leq n$. Notice that these points cannot be covered by the S-neighbor of u , since that block has a S-boundary point and its horizontal offset is positive. A normal connector between z and u would ensure that the vertical offset of u is at most as large as the horizontal offset of z . This modified connector forces a stronger requirement: it ensures that the vertical offset of u is at most the horizontal offset of z minus 1. A similar connector forces the horizontal offset of r to be at most the vertical offset of w minus 1.

The construction described so far depends only on k and n , but not on the structure of the graph G . The *tester points* in each tester gadget encode the edges of G . Let (x, y) be the coordinates of the common corner of blocks a, b, c, d in gadget $T_{j_1, (j_1, j_2)}$ connected to row j_1 and column (j_1, j_2) . If the p -th vertex is *not* the smaller endpoint of the q -th edge ($1 \leq p, q \leq n$), then we add the point $(x - q\epsilon, y - p\epsilon)$. The gadget $T_{j_2, (j_1, j_2)}$ is similarly defined, but in this case a tester point $(x - q\epsilon, y - p\epsilon)$ signifies that the p -th vertex is not the larger end point of the q -th edge. This completes the description of the reduction.

To prove the correctness of the reduction, first we show that if there is a solution for COVERING POINTS WITH SQUARES, then there is a size k clique in G . Recall that the problem parameter (the maximum number of allowed squares) in the constructed instance of COVERING POINTS WITH SQUARES equals the number of blocks. This means that the squares in the solution correspond to the blocks. As we have seen above, each row has a horizontal offset between 1 and n , let v_i be the vertex indexed by the horizontal offset of row i . Similarly, let e_{j_1, j_2} be the edge indexed by the vertical offset of column (j_1, j_2) .

We claim that the v_i 's form a clique in G , with e_{j_1, j_2} being the edge connecting v_{j_1} and v_{j_2} . Suppose that this claim does not hold for some j_1 and j_2 . Assume without loss of generality that e_{j_1, j_2} is not incident to v_{j_1} . Let us look at what happens in tester gadget $T_{j_1, (j_1, j_2)}$. Let p be the index of v_{j_1} and q be the index of e_{j_1, j_2} . The horizontal offset of row j_1 is p , hence the horizontal offset of t is at least p , and the vertical offset of u is at most $p - 1$. If we follow the implications of this, then it turns out that the vertical offset of a is at most $-p$, and the vertical offset of d is at least $-p + 1$. Similarly, the horizontal offset of c is forced to be at most $-q$, and the horizontal offset of b is forced to be at least $-q + 1$. Let (x, y) be the coordinates of the common corner of a, b, c , and d . From the assumption that the q -th edge is not incident to the p -th vertex, it follows that there is a tester point at $(x - q\epsilon, y - p\epsilon)$. This point is not covered by any of the blocks a, b, c, d : for example, block a cannot cover a point with vertical coordinate at least $y - p\epsilon$; block b cannot cover a point with horizontal coordinate less than $x - q\epsilon + \epsilon$, etc. This gives a contradiction.

To prove the other direction, we have to show that if there is a size k clique in G , then there is a solution for the constructed instance of COVERING POINTS WITH SQUARES. Let v_i be the i -th vertex in the clique, and let e_{j_1, j_2} be the edge connecting v_{j_1} and v_{j_2} . It is clear that we can cover the points in the k rows, $\binom{k}{2}$

columns, horizontal bands, and vertical bands in such a way that the horizontal offset of row i is the index of v_i , and the vertical offset of column (j_1, j_2) is the index of e_{j_1, j_2} . The only thing that should be verified is whether the points in the tester points in the tester gadgets can be covered. Consider the tester gadget that is connected to row j_1 and column (j_1, j_2) . Let p be the index of v_{j_1} , and let q be the index of e_{j_1, j_2} . Each tester point has coordinates $(x - \alpha\epsilon, y - \beta\epsilon)$ for some α and β . We use block c to cover all the tester points with $\alpha < q$; we use block b to cover the tester points with $\alpha \geq q + 1$, etc. By construction, there is no tester point at $(x - q\epsilon, y - p\epsilon)$, hence all the points are covered. \square

References

1. P. K. Agarwal, M. van Kreveld, and S. Suri. Label placement by maximum independent set in rectangles. *Comput. Geom.*, 11(3-4):209–218, 1998.
2. J. Alber and J. Fiala. Geometric separation and exact solutions for the parameterized independent set problem on disk graphs. *J. Algorithms*, 52(2):134–151, 2004.
3. S. Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *FOCS 1996*, pages 2–11. IEEE Comput. Soc. Press, 1996.
4. S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.
5. C. Bazgan. Schémas d’approximation et complexité paramétrée. Technical report, Université Paris Sud, 1995.
6. M. Cesati and L. Trevisan. On the efficiency of polynomial time approximation schemes. *Inform. Process. Lett.*, 64(4):165–171, 1997.
7. R. G. Downey. Parameterized complexity for the skeptic. In *Proceedings of the 18th IEEE Annual Conference on Computational Complexity*, pages 147–169, 2003.
8. R. G. Downey and M. R. Fellows. *Parameterized complexity*. Monographs in Computer Science. Springer-Verlag, New York, 1999.
9. T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric graphs. In *SODA 2001*, pages 671–679. SIAM, 2001.
10. D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985.
11. H. B. Hunt, III, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *J. Algorithms*, 26(2):238–274, 1998.
12. E. Malesińska. *Graph-Theoretical Models for Frequency Assignment Problems*. PhD thesis, Technical University of Berlin, 1997.
13. L. Sunil Chandran and F. Grandoni. Refined memorization for vertex cover. *Inform. Process. Lett.*, 93(3):125–131, 2005.
14. D. W. Wang and Y.-S. Kuo. A study on two geometric location problems. *Inform. Process. Lett.*, 28(6):281–286, 1988.