Grid Obstacle Representations With Connections to Staircase-Guarding

Therese Biedl[1] and Saeed Mehrabi[1]

[1] Cheriton School of Computer Science
University of Waterloo, Waterloo, Canada.
{biedl, smehrabi}@uwaterloo.ca

Abstract. In this paper, we study grid-obstacle representations of graphs where we assign grid-points to vertices and define obstacles such that an edge exists if and only if an xy-monotone grid path connects the two endpoints without hitting an obstacle or another vertex. It was previously argued that all planar graphs have a grid-obstacle representation in 2D, and all graphs have a grid-obstacle representation in 3D. In this paper, we show that such constructions are possible with significantly smaller grid-size than previously achieved. Then we study the variant where vertices are not blocking, and show that then grid-obstacle representations exist for bipartite graphs. The latter has applications in so-called *staircase-guarding* of orthogonal polygons; using our grid-obstacle representations we can argue that staircase-guarding is NP-hard in 2D.

1 Introduction

Recently, Bishnu et al. [7] initiated the study of *grid-obstacle representations*. Here the vertices of a graph $G = (V, E)$ are mapped to points in a grid, and other grid-points are marked as *obstacles* in such a way that (v, w) is an edge of G if and only if there exists an xy-monotone path in the grid from v to w that contains no obstacle-point and no point that belongs to some vertex $\neq v, w$. See also Fig. 1b. This is a special case (using the L_1-distance) of a more general problem, which asks for placing points and obstacles in the plane such that an edge (v, w) exists if and only if there is a shortest path (in a fixed distance metric) from v to w that does not intersect obstacles. See also Alpert et al. [2], who initiated the study of obstacle numbers, and [10] and the references therein for more recent developments.

Bishnu et al. [7] showed that any planar graph has a grid-obstacle representation in 2D, and every graph has a grid-obstacle representation in 3D. The main idea was to use a straight-line drawing, and then approximate it by putting a sufficiently fine grid around it that consists of obstacles everywhere except near the edge. The analysis of how fine a grid is required is not straightforward; Bishnu et al. claimed that in 2D an $O(n^2) \times O(n^2)$-grid is sufficient. They did not give bounds for the size needed in 3D (but it clearly is polynomial and at least $\Omega(n^2)$ in each dimension). Pach showed that not all bipartite graphs have grid-obstacle representations [14].

In this paper, we improve the grid-size bounds of [7]. In particular, rather than converting a straight-line drawing directly into a grid-obstacle representation, we first convert it into a visibility representation or an orthogonal drawing that has special properties, but resides in a linear-size grid. This can then be easily converted to a grid-obstacle-representation. Thus we obtain 2D grid-obstacle representations for planar graphs in an $O(n) \times O(n)$-grid, and 3D grid-obstacle representations for all graphs in an $O(n) \times O(n) \times O(n)$-grid.

We then discuss the case with the restriction that vertices act as obstacles for edges not incident to them, and show that sometimes this restriction can be dropped. We hence obtain *non-blocking grid-representations* in 2D for all planar bipartite graphs and in 3D for arbitrary bipartite graphs.

The latter has applications: we can use the constructions for hardness proofs for a polygon-guarding problem. A point guard g is said to *staircase-guard* (or *s-guard* for short) a point p inside an orthogonal polygon P if p can be reached from g by a *staircase*; that is, an orthogonal path inside P that is both x- and y-monotone. In the *s-guarding problem*, the objective is to guard an orthogonal polygon with the minimum number of s-guards. Motwani et al. [13] proved that s-guarding is polynomial on simple orthogonal polygons. Gewali and Ntafos [12] proved that the problem is NP-hard in 3D; since they reduce from vertex cover in graphs with maximum degree 3 this in fact implies APX-hardness in 3D [1]. But to our knowledge, the complexity was open for 2D polygons with holes. Using non-blocking grid-representations, we show that it is NP-hard.

2 2D Grid Obstacle Representations

Let $G = (V, E)$ be a planar graph. To build a grid-obstacle representation, we use a *visibility representation* where every vertex is represented by a *bar* (a horizontal line segment), and every edge is represented by a vertical line segment between its ends [17, 16, 15]. We need here a construction with a special property, which can easily be achieved by "shifting around" where the edges attach at the vertices. (A direct proof of this is given in the appendix.) See also Fig. 1a.

Lemma 1. *Every planar graph has a visibility representation in an $O(n)\times O(n)$-grid for which any vertex-bar can be split into a left and right part such that all downward edges attach on the left and all upward edges attach on the right.*

Now convert such a visibility representation into a grid-obstacle representation. First, double the grid so that no two grid-points on edge-segments or vertex-bars are adjacent unless the corresponding graph-elements were. For each vertex v, assign as vertex-point some grid-point that lies between the two parts of the bar of v; this exists since we doubled the grid. The obstacles consist of all grid points that are not on some edge segment or vertex bar. Clearly, the representation is in an $O(n) \times O(n)$-grid. One can easily show (see the appendix) that this is a grid-obstacle representation, and so we have:

Theorem 1. *Every planar graph has a 2D grid-obstacle representation in an $O(n) \times O(n)$-grid.*

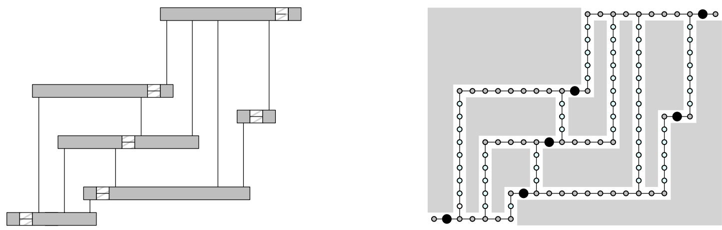

(a) (b)

Fig. 1: A special visibility representation gives a grid-obstacle representation.

One can easily argue that any straight-line drawing of a planar graph of height H can be converted into a visibility representation of height $2H$ and width $O(n)$ (see also [3]). Then we can apply the same approach as above. Based on drawings for trees [9], outer-planar graphs [4] and series-parallel graphs [4], we hence get:

Corollary 1. *Every tree and every outer-planar graph has a 2D grid-obstacle representation in an* $O(\log n) \times O(n)$*-grid. Every series-parallel graph has a 2D grid-obstacle representation in an* $O(\sqrt{n}) \times O(n)$*-grid.*

3 3D Obstacle Representation

In this section, we argue that a similar (and even simpler) construction gives a grid-obstacle representation in 3D. We obtain this by building an orthogonal representation first that has special properties. This representation is not quite a graph drawing, because edges may overlap; this will not create problems for the obstacle representation later.

Enumerate the vertices as $v_1, \ldots, v_n$ in arbitrary order. Place v_i at (i, i, i). To draw an edge (v_i, v_j) with $i < j$, we use the path $(i, i, i)–(j, i, i)–(j, i, j)–(j, j, j)$ along the cube spanned between the two points. See Fig. 2. Observe that all edges (v_h, v_i) with $h < i$ reach v_i from the y^--side and that all edges (v_i, v_j) with $i < j$ leave v_i at the x^+-side. Edges incident to v_i may overlap along these two sides, but otherwise there are no overlaps or crossings in the drawing. Also, we clearly reside in an $n \times n \times n$-grid.

Now double the grid, then cover any grid-point by an obstacle unless it is used by a vertex or an edge. One can easily argue that the result is a grid-obstacle representation (see the appendix), and we have:

Theorem 2. *Every graph has a 3D grid-obstacle representation in an* $O(n) \times O(n) \times O(n)$*-grid.*

Notice that the obstacle in this case can be made to be just one polyhedron (albeit of high genus).

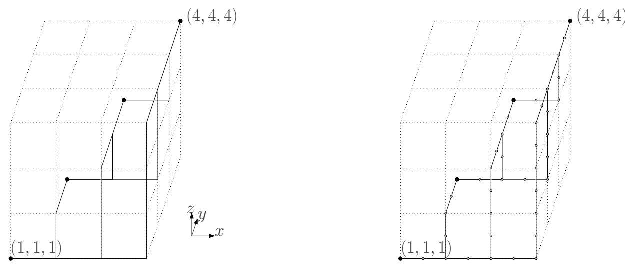

Fig. 2: A 3D orthogonal representation of K_4, and turning it into a grid-obstacle representation. All grid-points that are not shown are blocked by obstacles.

4 Non-blocking grid-obstacle representations

In our definition of grid-obstacle representation, we required that the grid point of any vertex v acts as an obstacle to any other path. The main reason for this is that otherwise paths could "seep through" a vertex, creating unwanted adjacencies. In this section, we consider *non-blocking grid-obstacle representations*, which means that vertices do not act as obstacles.

4.1 Planar bipartite graphs

We first give an algorithm for non-blocking grid-obstacle representation of planar bipartite graphs. It is known that any such graph $G = (A \cup B, E)$ has an *HH-drawing* [6], i.e., a planar drawing where all vertices in A have positive y-coordinate, all vertices in B have negative y-coordinate, every edge is drawn with at most one bend, and all bends have y-coordinate 0. See also Fig. 3. In particular, we know that every edge is drawn y-monotonically. Any such a drawing can be converted into a visibility representation [5] where the y-coordinate of every vertex is unchanged. So, we obtain:

Lemma 2. *Let $G = (A \cup B, E)$ be a planar bipartite graph. Then, there exists a visibility representation of G such that all vertices in A have only neighbours below, and all vertices in B have only neighbours above.*

Now create an obstacle representation as before by doubling the grid, and placing obstacles at all grid-points that are not used by the drawing. Place each vertex $a \in A$ at the rightmost grid-point of the bar of a, and each $b \in B$ at the leftmost grid-point of the bar of b. One easily verifies that this is a non-blocking grid-obstacle representation: For each vertex a in A, no xy-monotone path can go through the grid-point of a without ending there, because no grid-point higher than a can be reached when going through a. Similarly one argues for B, and so we have:

Theorem 3. *Every planar bipartite graph has a non-blocking grid-obstacle representation in an $O(n) \times O(n)$-grid.*

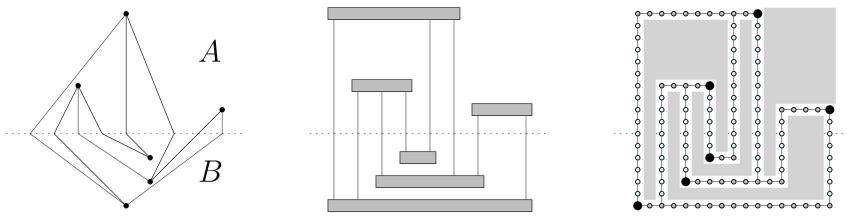

Fig. 3: An HH-drawing of a planar bipartite graph, and converting it to a non-blocking grid-obstacle representation.

4.2 Application to *s*-guarding

Recall that the s-guarding problem consists of finding the minimum set S of points in a given orthogonal polygon P such that for any $q \in P$ there exists a $p \in S$ that is connected to q via a staircase inside P. Using non-blocking grid-obstacle representations, we can show:

Theorem 4. *s-guarding is* NP*-hard on orthogonal polygons with holes.*

Proof. We reduce from minimum dominating set, i.e., the problem of finding a set D of vertices in a graph such that every vertex is either in D or has a neighbor in D. This is NP-hard, even on planar bipartite graphs [8]. Given a planar bipartite graph $G = (A \cup B, E)$, construct the non-blocking grid-obstacle representation Γ. Let P' consist of all unit squares (*pixels*) around grid-points that are *not* in an obstacle. The obstacles of Γ become holes in P'. Now for any vertex $a \in A$ extend the bar of a slightly rightward beyond the last edge, and for every $b \in B$ extend the bar leftward beyond the last edge. Finally, at every edge e, attach two "spirals" on the left and right side of its vertical segment; the one on the left curls upward while the on the right curls downward. See Fig. 4. These spirals are small enough that they fit within the holes of P', without overlapping other parts of P' or each other. Call the resulting polygon P. We can show (see the appendix) that G has a dominating set of size k if and only if P can be s-guarded with $2|E| + k$ guards. This proves the theorem. □

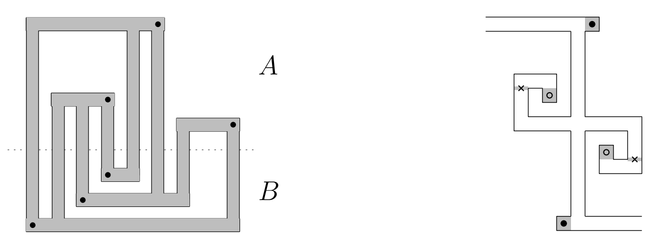

(a) (b)

Fig. 4: The polygon for the graph in Fig. 3, and gadgets that we attach.

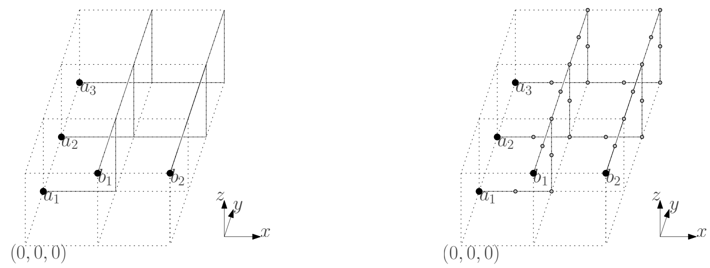

Fig. 5: A 3D orthogonal representation of $K_{2,3}$, and turning it into a grid-obstacle representation. Grid-points not shown are covered by obstacles.

4.3 3D Obstacle Representation of bipartite graphs

In 3D, all bipartite graphs have a non-blocking grid-obstacle representation: Enumerate the vertices as $A = \{a_1, \dots, a_\ell\}$ and $B = \{b_1, \dots, b_k\}$. Place a point for vertex a_i at $(0, i, 0)$ and a point for vertex b_j at $(j, 0, 1)$. Route each edge (a_i, b_j) as the orthogonal path $(0, i, 0) - (j, i, 0) - (j, i, 1) - (j, 0, 1)$, and observe that two paths overlap in the x^+-direction at a_i if they both begin at a_i, or overlap in the y^+-direction at b_j if they both end at b_j, but otherwise there is no overlap. Now obtain the grid-obstacle representation as before by doubling the grid and making grid-points obstacles unless they are used by vertices and edge-paths. As before one argues that this is indeed a non-blocking grid-obstacle-representation and so we have:

Theorem 5. *Every bipartite graph has a 3D non-blocking grid-obstacle representation.*

5 Conclusion

In this paper, we studied grid-obstacle representations. We gave constructions with smaller grid-size for planar graphs in 2D and all graphs in 3D. If the graph is bipartite then we can construct representations where vertices are not considered obstacles. We used these types of representation to prove NP-hardness of the s-guarding problem in 2D polygons with holes.

The main interesting open question is whether fewer obstacles might be enough. Currently, even if we allow obstacles to be polygons rather than grid-points, we would need (in Theorems 1 and 3) one obstacle per face of the planar graph, or $\Theta(n)$ in total. For obstacle-representations that us straight-line segments, rather than xy-monotone grid-paths, significantly fewer obstacles suffice [10]. Can we create grid-obstacle representations with $o(n)$ obstacles, at least for some subclasses of planar graphs? Another direction for future work would be to find other classes of graphs for which we can construct non-blocking grid-obstacle representation. Does this exist for all planar graphs in 2D?

References

1. Paola Alimonti and Viggo Kann. Some apx-completeness results for cubic graphs. *Theor. Comput. Sci.*, 237(1–2):123–134, 2000.
2. Hannah Alpert, Christina Koch, and Joshua D. Laison. Obstacle numbers of graphs. *Discrete & Computational Geometry*, 44(1):223–244, 2010.
3. T. Biedl. Height-preserving transformations of planar graph drawings. In C. Duncan and A. Symvonis, editors, *Graph Drawing (GD'14)*, volume 8871 of *LNCS*, pages 380–391. Springer, 2014.
4. Therese Biedl. Small drawings of outerplanar graphs, series-parallel graphs, and other planar graphs. *Discrete & Computational Geometry*, 45(1):141–160, 2011.
5. Therese Biedl. Height-preserving transformations of planar graph drawings. In *proceedings of the 22nd International Symposium on Graph Drawing (GD '14), Würzburg, Germany, Revised Selected Papers*, pages 380–391, 2014.
6. Therese Biedl, Michael Kaufmann, and Petra Mutzel. Drawing planar partitions II: HH-drawings. In *proceedings of the 24th International Workshop on Graph-Theoretic Concepts in Computer Science (WG '98)*, pages 124–136, 1998.
7. A. Bishnu, A. Ghosh, R. Mathew, G. Mishra, and S. Paul. Grid obstacle representations of graphs, 2016. Manuscript.
8. Brent N. Clark, Charles J. Colbourn, and David S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1-3):165–177, 1990.
9. Pierluigi Crescenzi, Giuseppe Di Battista, and Adolfo Piperno. A note on optimal area algorithms for upward drawings of binary trees. *Comput. Geom.*, 2:187–200, 1992.
10. Vida Dujmovic and Pat Morin. On obstacle numbers. *Electr. J. Comb.*, 22(3):P3.1, 2015.
11. H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10:41–51, 1990.
12. Laxmi Gewali and Simeon C. Ntafos. Covering grids and orthogonal polygons with periscope guards. *Comput. Geom.*, 2:309–334, 1992.
13. Rajeev Motwani, Arvind Raghunathan, and Huzur Saran. Covering orthogonal polygons with star polygons: The perfect graph approach. *J. Comput. Syst. Sci.*, 40(1):19–48, 1990.
14. J. Pach. Graphs with no grid obstacle representation. *Geombinatorics*, 26(2):80–83, 2016.
15. P. Rosenstiehl and R. E. Tarjan. Rectilinear planar layouts and bipolar orientation of planar graphs. *Discrete Computational Geometry*, 1:343–353, 1986.
16. R. Tamassia and I. Tollis. A unified approach to visibility representations of planar graphs. *Discrete Computational Geometry*, 1:321–341, 1986.
17. S. Wismath. Characterizing bar line-of-sight graphs. In *ACM Symposium on Computational Geometry (SoCG '85)*, pages 147–152, 1985.

A Missing Proofs

Proof of Lemma 1 Take a planar straight-line drawing of G that has height $O(n)$ and where no edge is drawn horizontally. (For example, the drawing of de Fraysseix et al. [11] is easily seen to achieve this if we modify the placement of the initial triangle v_1, v_2, v_3.) Direct the edges from the lower endpoint to the higher endpoint. Now split each vertex v into two adjacent vertices v^ℓ and v^r, where v^ℓ is adjacent to all incoming edges of v and v^r is adjacent to all outgoing edges. Double the height of the drawing by inserting a new row after each existing one. Place v^r to the right of v, and re-route all outgoing edges of v to leave from v^r instead, by adding a bend in the row above v. See Fig. 6.

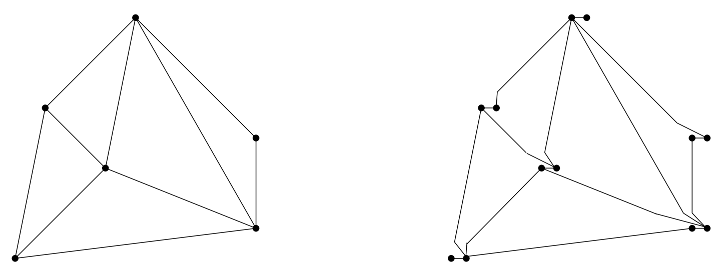

(a) (b)

Fig. 6: Split the vertices of a planar drawing. Converting this into a visibility representation gives the one shown in Fig. 1.

Observe that all edges are drawn y-monotonically, and (v^ℓ, v^r) is drawn horizontally for all vertices v. Now we convert this drawing, using the method of [5], into a visibility representation such that the y-coordinates and the order of vertices within a row is unchanged. This has again height $O(n)$, and (after deleting empty columns) it also has width $O(m + n) = O(n)$ [5]. The two vertices v^ℓ and v^r that replaced vertex v were consecutive in one row, hence in the visibility representation they are also consecutive (and the edge between them is horizontal) and we can simply recombine them to obtain one horizontal segment for vertex v. This gives the required visibility representation.

Proof of Theorem 1 We must argue that the constructed representation is a grid-obstacle representation. For any vertex v, let v^ℓ and v^r be the left and the right part of the bar of v where downward/upward edges attach. Consider an edge (v_i, v_k) where v_i has the smaller y-coordinate. In the visibility representation, the corresponding segment attaches at v_i^r and v_k^ℓ. In the obstacle representation, we can hence go rightward from the grid-point of v_i along edge (v_i^ℓ, v_i^r) and the segment of v_i^r, then up along the segment of (v_i^r, v_k^ℓ), then rightward along v_k^ℓ and the edge (v_k^ℓ, v_k^r) to reach v_k along an xy-monotone grid-path.

On the other hand, assume that v_i and v_k have an xy-monotone grid-path π between them in the obstacle-representation. No two vertex-points have the same y-coordinate, so we may assume that the point of v_i is lower. If π left v_i on the left side, then it would have to continue downward from there, which contradicts monotonicity. So π leaves v_i on the right side. From there, it can only go upward along some edge (v_i, v_j) and reach the segment of v_j^ℓ. All edges attaching here go downward, which π cannot use by monotonicity. So π must continue to the grid-point of v_j. Here π is obstructed if $v_j \neq v_k$, so we must have $v_j = v_k$ and (v_i, v_k) is an edge as desired.

Proof of Theorem 2 We must argue that the constructed representation is a grid-obstacle representation. Clearly, for any edge (v_i, v_k) we can find an xy-monotone path by walking along the route of (v_i, v_k). Vice versa, if there is a monotone path π from v_i to v_k with (say) $i < k$, then it must connect (i, i, i) to (k, k, k) and so be going in positive direction. Thus it must leave v_i on the x^+-side. From here the only option is to continue in z^+-direction starting at some point (i, j, i). This necessarily leads to (i, j, j), since there are no other adjacent unobstructed grid-points. From there the only positive direction possible is to go to (j, j, j). But then v_j blocks the path, so we must have $v_j = v_k$ and (v_i, v_k) is an edge as desired.

Proof of Theorem 4 We aim to show that G has a dominating set of size k if and only if P can be guarded by $k + 2|E|$ s-guards. Recall that in P every vertex corresponds to a bar (of the visibility representation of G) and every edge e corresponds to a channel (along the vertical segment that represented e). Also, for each vertex u we attached an "end-pixel" ψ_u that is beyond all attachment points of all edge-channels; this is on the right end of the bar if $u \in A$ and on the left end if $u \in B$. These are marked by black dots in Fig. 4(b)).

For any edge (a, b), we also attached two spiral-gadgets to the edge-channel of (a, b). In any such spiral σ, there are two crucial places. One is the "tail-pixel" ψ_σ at the end of the spiral (marked by a circle in Fig. 4(b)), and the other is the line segment s_σ that marks the boundary of points that can s-guard ψ_σ (marked by a cross in Fig. 4(b)).

Consider a dominating set D of size k in G and define a set S of points in P as follows. For each $u \in D$, add an arbitrary point $p(u)$ of the end-pixel ψ_u to S. Observe that $p(u)$ guards ψ_u as well as ψ_v of any vertex v for which (u, v) is an edge. Secondly, for each edge-spiral σ, add the point x_σ marked with a cross in Fig. 4. Note that the point x_σ of the left spiral σ at an edge (a, b) can s-guard all of σ, the bottom half of the edge-channel for (a, b) and everything of the vertex-bar of b (the lower endpoint) that is to the right of where the edge-channel attaches. Similarly, the point in the right spiral can see the top half of the edge-channel and everything of the vertex-bar of a to the left of where the edge-channel attaches. All the spiral-guards together hence cover everything except the end-pixels, but those are guarded by the points added due to dominating set D. So the chosen $k + 2|E|$ points guard everything.

Conversely, let S be a set of $k + 2|E|$ points in P that s-guard all of P. For each spiral-gadget σ, there must exist some guard $s \in S$ that guards the tail-pixel. This guard must lie on segment s_σ (or even closer to the tail-pixel), and as one easily verifies, cannot see any point in an end-pixel, or any tail-pixel of any other spiral-gadget. Therefore, there are $2|E|$ such points in total (call them M), none of which guards an end-pixel.

This leaves at most k guards that s-guard all end-pixels. Define D as follows. For any vertex v, if some point in S lies in the vertex-bar of v, add v to D. For any edge $e = (u, v)$, if some point in S s-guards some end-pixel and lies on the edge-channel or a spiral-gadget of e, then arbitrarily add one of u and v to D. Since none of the guards in M fit this description, we have $|D| \leq k$. For any vertex v, the end-pixel $\psi(v)$ must have been guarded by some point $g \in S$. This implies that $p(v)$ s-guards g, which is possible only if g lies in the bar of v, the bar of some neighbour u of v, or the channel (or adjacent spirals) of some edge (v, u). Hence g gives rise to a vertex in D that is either v or a neighbour of v. Thus D is a dominating set as required.