Laboratoire Bordelais de Recherche en Informatique
UMR 5800 - Université Bordeaux I, 351, cours de la Libération,
33405 Talence CEDEX, France

Research Report RR-1258-01

Small k-Dominating Sets in Planar Graphs with Applications

by Cyril Gavoille, David Peleg,
André Raspaud and Éric Sopena

May, 2001

Small k-Dominating Sets in Planar Graphs with Applications

Cyril Gavoille* David Peleg† André Raspaud*

Éric Sopena*

Abstract

A subset of nodes S in a graph G is called k-*dominating* if, for every node u of the graph, the distance from u to S is at most k. We consider the parameter $\gamma_k(G)$ defined as the cardinality of the smallest k-dominating set of G. For planar graphs, we show that for every $\epsilon > 0$ and for every $k \geqslant (5/7 + \epsilon)D$, $\gamma_k(G) = O(1/\epsilon)$. For several subclasses of planar graphs of diameter D, we show that $\gamma_k(G)$ is bounded by a constant for $k \geqslant D/2$. We conjecture that the same result holds for every planar graph. This problem is motivated by the design of routing schemes with compact data structures.

Keywords: dominating sets, planar graphs, compact routing tables

*LaBRI, Université Bordeaux I, 351, cours de la Libération, 33405 Talence Cedex, France. {*gavoille,raspaud,sopena*}@*labri.fr*

†Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. *peleg@wisdom.weizmann.ac.il*

1 Introduction

Given a graph $G = (V, E)$, denote by $d_G(u, v)$ the distance in G from u to v, and by $\Gamma_k(u) = \{v \in V \mid d_G(u, v) \leqslant k\}$ the ball of radius k centered at u (note that $\Gamma_k(u)$ is also defined for non integral k). We extend this notation to every subset $S \subseteq V$, letting $\Gamma_k(S) = \bigcup_{u \in S} \Gamma_k(u)$. We say that a subset of nodes $S \subseteq V$ is a *k-dominating set* for G (or, S k-dominates G, or S covers G) if $\Gamma_k(S) = V$. Let $\gamma_k(G)$ denote the cardinality of the smallest k-dominating set of G. Our goal is to bound $\gamma_k(G)$ for various planar graphs G and values of k.

To illustrate the properties of the parameter $\gamma_k(G)$, let us consider a tree T of diameter D. It is clear that if $k \geqslant D/2$, then $\gamma_k(T) = 1$; it suffices to consider the center of T. On the other hand, if $k < D/2$ then there exists some n-node tree T_0 of diameter D for which $\gamma_k(T_0) \geqslant 2(n-1)/D$. For instance, consider T_0 composed of a star $K_{1,p}$ with each edge subdivided into $q-1$ nodes, where $n = pq+1$, $D = 2q$ and $\gamma_{q-1}(T_0) = p$. Every two leaves of this tree are at distance $2q$, thus p nodes of T_0 are required to $(q-1)$-dominate all the leaves.

More generally, for $k \geqslant D/2$, every n-node graph $G = (V, E)$ of diameter D satisfies $\gamma_k(G) < \sqrt{n(1 + \ln n)}$. To see this, we use a dual characterization of a k-dominating set as a set S which *hits* the collection $\{\Gamma_k(u) \mid u \in V\}$, i.e., such that $S \cap \Gamma_k(u) \neq \varnothing$ for every $u \in V$. Now, note that for $k \geqslant D/2$, $\Gamma_k(x) \cap \Gamma_k(y) \neq \varnothing$ for any two nodes x and y. Thus every set $\Gamma_k(x)$ k-dominates G. Hence, either there exists a node x such that $|\Gamma_k(x)| < \alpha$, where $\alpha = \sqrt{n(1 + \ln n)}$, and we are done, or $|\Gamma_k(x)| \geqslant \alpha$ for every node x. In the latter case, the claim follows by a result of Lovász [Lov75] about cover sets, which states that there exists a k-dominating set of G of size $\beta < n(1 + \ln n)/\min_x |\Gamma_k(x)| \leqslant \sqrt{n(1 + \ln n)}$.

For planar graphs, it is known that the size of a dominating set (i.e., 1-dominating set) is bounded by 3 if $D = 2$, and by 10 if $D = 3$, cf. [MS96]. Thus, for $k \geqslant D/2$, and every planar graph G, $\gamma_k(G) \leqslant 3$ if $D = 2$, and $\gamma_k(G) \leqslant 10$ if $D = 3$.

A recent result concerns planar triangulations. A triangulation of the plane is a planar graph with an embedding on the plane such that each face, except maybe the outer face, is a triangle. In [CV00] it is shown that every planar triangulation G in which every internal node (i.e., that does not belong to the outer face) has degree at least 6, satisfies $\gamma_k(G) \leqslant 2$, again for $k \geqslant D/2$.

Section 2 presents our main results. For planar graphs, we show that for every $\epsilon > 0$ and for every integer $k \geqslant (5/7+\epsilon)D$, $\gamma_k(G) = O(1/\epsilon)$. For outerplanar graphs of diameter D, we show that $\gamma_k(G) \leqslant 2$ for $k \geqslant D/2$. We conjecture that $\gamma_k(G)$ is bounded by a constant for every planar graph of diameter D for $k \geqslant D/2$.

In Section 3 we show study bounded treewidth and chordality graphs. More precisely, we show that if G has treewidth bounded by t, then $\gamma_k(G) \leqslant t+1$ for $k \geqslant \lfloor D/2 \rfloor$, and $\gamma_k(G) = 1$ for $k \geqslant \lfloor D/2 \rfloor + \lfloor c/2 \rfloor$ if G has no chordless cycle of length greater than c.

The motivation for studying this parameter stems from the design of routing schemes with compact data structures. In particular, it is shown in Section 4 that for every graph G of diameter D and every $k \geqslant 0$, G has an interval routing scheme [SK85, vLT87] with dilation at most $D+k$ and compactness at most $\frac{1}{2}\gamma_k(G)+1$. (The dilation of a routing scheme is the length of the longest route, and its compactness is the maximum number of intervals over any arc.) Moreover, the upper bound on the compactness can be reduced to $\frac{1}{4}\gamma_k(G)+o(\gamma_k(G))$, if $\gamma_k(G)/\log n \to +\infty$. This result improves (by a multiplicative factor) and generalizes the compactness vs. dilation trade-offs of [KRŠ00] and of [TL99]. Our result also implies that for every constant $\epsilon > 0$, every planar graph of diameter D has an interval routing scheme with dilation at most $\lceil (12/7+\epsilon)D \rceil$ and constant compactness.

2 Plane Graphs

2.1 Preliminaries

Hereafter, we assume that $G = (V, E)$ is connected and of diameter D. The following basic property is important for the remainder of the paper.

Proposition 2.1 *Let $S \subset V$, let $k \geqslant \lfloor D/2 \rfloor$, and let $u, v \notin \Gamma_k(S)$. Then no shortest path between u and v goes through any node of S.*

Proof. Assume, towards contradiction, that there is some shortest path from u to v that contains a node $w \in S$. Then $d_G(u,v) = d_G(u,w) + d_G(w,v)$.

Since $u, v \notin \Gamma_k(S)$, necessarily $u, v \notin \Gamma_k(w)$. Thus $d_G(u,w) > k$, and more precisely, $d_G(u,w) \geqslant \lfloor k \rfloor + 1$ since the distance is integral. Similarly $d_G(v,w) \geqslant \lfloor k \rfloor + 1$. It implies that $d_G(u,v) \geqslant 2\lfloor k \rfloor + 2 \geqslant 2\lfloor D/2 \rfloor + 2 \geqslant D+1$, in contradiction with the fact that D is the diameter of G. ∎

A subset $S \subset V(G)$ is a *separator* of G if $G \setminus S$ is composed of two or more connected components.

Proposition 2.1 yields the following immediate corollary, keeping in mind that every path joining two nodes of different connected components in $G \setminus S$, for some separator S, has to go through S.

Corollary 2.2 *Let $k \geqslant \lfloor D/2 \rfloor$, let S be a separator of G, and let $U = V \setminus \Gamma_k(S)$. Then there exists a unique connected component C of $G \setminus S$ such that $U \subseteq V(C)$.*

2.2 Outerplanar Graphs

Let us start by proving that for outerplanar graphs, $\gamma_k(G)$ is bounded by a constant for $k \geqslant D/2$. Recall that a graph H is a *minor* of a graph G if H can be obtained from G by a sequence of zero or more node deletions, edge deletions or edge contractions. We say that G is *H-free* if it does not contain H as a minor. Let us denote by $K_{p,q}$ the complete bipartite graph with p nodes in one partition and q nodes in the other. We prove the following.

Theorem 2.3 *If G is $K_{2,t+1}$-free then $\gamma_k(G) \leqslant t$ for every $k \geqslant D/2$, and this is tight for a subdivision of $K_{t,t}$.*

Proof. Let u and v be any two nodes in G with $d_G(u,v) = D$. The *level* of the node x is defined as its distance from u, denoted $l(x) = d_G(u,x)$. Let $V_i = \{x \in V \mid l(x) = i\}$ for every $0 \leqslant i \leqslant D$. Let M be the set of nodes $x \in V_k$ such that there exists a path $(x, x_1, x_2, \ldots, x_\ell, v)$ from x to v satisfying $l(x_i) > k$ for every $1 \leqslant i \leqslant \ell$.

We first claim that the set M has cardinality at most t. To prove this, we show that if M contains $t+1$ distinct nodes then G contains $K_{2,t+1}$ as a minor, a contradiction. Indeed, consider the graph H composed of the $t+1$ nodes of M and the $2(t+1)$ paths linking M to u and v, and delete the edges and nodes of $G \setminus H$. Then it is easy to see that by contractions we obtain a copy of $K_{2,t+1}$.

Now let $P = (u, u_1, u_2, \ldots, u_p, v)$ be any path linking u and v and let q be the largest subscript such that $l(u_q) = k$. Clearly, $l(u_i) > k$ for every i in the range $q < i \leqslant p$, and thus $u_q \in M$. Hence every path linking u and v must go through the set M which, therefore, is a separator in G. (In particular, u and v belong to different connected components of $G \setminus M$).

We now prove (by contradiction again) that M k-dominates G. Suppose that there exists a node x in G such that $d_G(x, M) > k$. As observed above, M separates u from v, hence either u or v does not belong to the connected component of x in $G \setminus M$. In the former case we get

$$d_G(u,x) \ \geqslant\ d_G(u,M) + d_G(x,M) \ >\ 2k \geqslant D,$$

and in the latter case

$$d_G(v,x) \ \geqslant\ d_G(v,M) + d_G(x,M) \ >\ D - k + k = D\ .$$

Therefore, either $d_G(u,x)$ or $d_G(v,x)$ is strictly greater than D, a contradiction.

We thus get that M k-dominates G, hence $\gamma_k(G) \leqslant |M| \leqslant t$, completing the proof. ∎

As every outerplanar graph is $K_{2,3}$-free, we get the following.

Corollary 2.4 *If G is an outerplanar graph of diameter D then $\gamma_k(G) \leqslant 2$, for every $k \geqslant D/2$.*

2.3 Planar Graphs

We assume from now on that G is a plane graph, that is, a planar graph with an embedding in $\mathbb{R}^2$. More precisely, the nodes are points of $\mathbb{R}^2$ and the edges are simple curves that can cross or meet each other on the nodes only. A connected subset of $\mathbb{R}^2$ is called a *region* of the plane. Given a subgraph H of G and a point $w \notin H$ (i.e., a point $w \in \mathbb{R}^2$ that does not belong to an edge or a node of H), we denote by $\operatorname{reg}(w, H)$ the unique region of $\mathbb{R}^2 \setminus H$ containing w. Note that if H is a tree, then $\mathbb{R}^2 \setminus H$ consists of one region only. If $\mathbb{R}^2 \setminus H$

is composed of exactly two non-empty regions, we denote by $\overline{\text{reg}}(w, H)$ the complementing region, such that $\mathbb{R}^2 = \text{reg}(w, H) \cup \overline{\text{reg}}(w, H) \cup H$.

Given an integer $\psi \geqslant 3$, a *ψ-gon* w.r.t. G and k is a subgraph H of G defined by a sequence $(x_1, \ldots, x_\psi)$ of distinct nodes, and a sequence $(Q_1, \ldots, Q_\psi)$ of paths such that the following four conditions hold:

1. Q_i is a shortest path from x_i to $x_{(i \bmod \psi)+1}$, for every $i \in \{1, \ldots, \psi\}$;
2. $\mathbb{R}^2 \setminus H$ is composed of at most two non-empty regions;
3. $x_\psi \notin \Gamma_k(Q_1 \cup \cdots \cup Q_{\psi-2})$;
4. $d(x_i, x_{(i \bmod \psi)+1}) > k$, for every $i \in \{1, \ldots, \psi\}$.

The point x_ψ is referred to as the *extremity* of the ψ-gon. A ψ-gon H is called *weak* if conditions 3 and 4 do not necessarily hold. A weak 3-gon is hereafter called a *triangle*.

Roughly speaking, a ψ-gon consists of ψ shortest paths, each of length at least $k+1$, possibly sharing some edges. Fig. 1 presents a 4-gon H w.r.t. a plane graph G of diameter $D = 4$ and $k = 2$. Note that the x_i's may border to different regions induced by H.

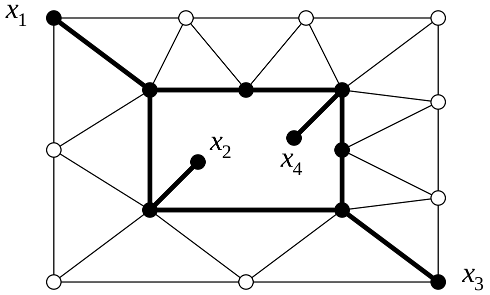

Figure 1: A plane graph G with a 4-gon H for $k = 2$ (bold edges).

Given a path P and two nodes u and v on P, we denote by $P[u, v]$ the subpath of P between u and v.

Let $Q_1, \ldots, Q_p$ be a set of shortest paths, and for each $i \in \{1, \ldots, p\}$ let $(z^i_1, \ldots, z^i_{t_i})$ be the sequence of nodes of $P \cap Q_i$ encountered in this order along a walk on P from u to v (possibly $t_i = 0$ if $P \cap Q_i = \varnothing$). The *simplified path of P on $Q_1 \cup \cdots \cup Q_p$* is the path defined by

$$\begin{aligned} S \;=\; & P[u, z^1_1] \cup Q_1[z^1_1, z^1_{t_1}] \cup P[z^1_{t_1}, z^2_1] \cup \cdots \cup Q_i[z^i_1, z^i_{t_i}] \cup P[z^i_{t_i}, z^{i+1}_1] \cup \cdots \\ & \cup Q_p[z^p_1, z^p_{t_p}] \cup P[z^p_{t_p}, v] \;. \end{aligned}$$

Roughly speaking, S uses shortcuts of P on Q_1, then shortcuts of P on Q_2, and so on. Note that if P is a shortest path, then S is also a shortest path between u and v. (If $t_i = 0$ for all i, then $S = P$.) Fig. 2 shows the simplified path S of P on $Q_1 \cup Q_2$. Note that S is unique, but other shortcuts are possible, for instance, $Q_1[u, z_2^1] \cup P[z_2^1, z_1^2] \cup Q_2[z_1^2, v]$.

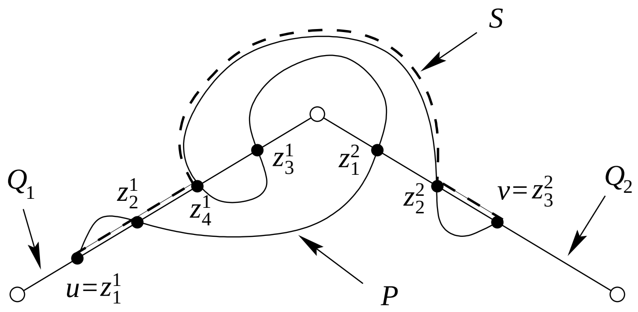

Figure 2: The simplified path S of P on $Q_1 \cup Q_2$.

For notational convenience, we use $\Gamma_k(H)$ where H is a graph instead of $\Gamma_k(V(H))$.

The next lemma shows an important property of ψ-gons.

Lemma 2.5 *Let H be a p-gon of G with $p \leqslant 4$ and let $k \geqslant \lfloor D/2 \rfloor$. Then for every $w \notin \Gamma_k(H)$:*

1. *there exists a shortest path A from w to x_p which is entirely contained in* $\mathrm{reg}(w, H) \cup H$*;*
2. *there exists a shortest path B from w to x_i, $i < p$, which is entirely contained in* $\mathrm{reg}(w, H) \cup H$.

Moreover, A and B can be chosen such that the two graphs $H_1 = B \cup (Q_i \cup Q_{i+1} \cup \cdots \cup Q_{p-1}) \cup A$ and $H_2 = B \cup (Q_{i-1} \cup Q_{i-2} \cup \cdots \cup Q_1) \cup Q_p \cup A$ are, respectively, a $(p-i+2)$-gon and an $(i+2)$-gon of G.

Proof. Note that to prove Lemma 2.5 it suffices to construct a shortest path A from w to x_p, and a shortest path B from w to x_i, $i < p$, that are wholly contained in $\mathrm{reg}(w, H) \cup H$. Indeed, to form the ψ-gons H_1 and H_2 we consider the simplified path S of B on A, and replace B by S. Then we check that H_1 and H_2 each induces at most two regions of $\mathbb{R}^2$, and the four conditions of the definition of ψ-gons follows.

Without loss of generality we assume that $\mathbb{R}^2 \setminus H$ is composed of two non-empty regions. Indeed, if $\overline{\mathrm{reg}}(w,H) = \varnothing$, then every shortest from w to x_i is wholly contains in $\mathrm{reg}(w,H) \cup H = \mathbb{R}^2$.

Let P be a shortest path from w to x_p. Let u be the node of $P \cap H$ closest to w, and assume that $u \in Q_\ell$. If $\ell \in \{p, p-1\}$, then $A = P[w,u] \cup Q_\ell[u,x_p] \subseteq \mathrm{reg}(w,H) \cup H$, and the first claim follows. So henceforth assume that $\ell \notin \{p,p-1\}$. Consider $S = \bigcup_{i=1}^{p-2} Q_i$. Note that Q_ℓ is contained in S, so $u \in P \cap S$. Now, $w \notin \Gamma_k(S)$ and $x_p \notin \Gamma_k(S)$ (by Condition 3 of the definition). By Proposition 2.1, the shortest path P between w and x_p cannot intersect S: a contradiction. Thus the path $A = P$ exists, completing the proof of first claim.

To prove the second claim, we proceed by induction on p.

Property 2.6 *Let T be a triangle defined by the node sequence (x_1,x_2,x_3) and the shortest paths (Q_1,Q_2,Q_3). For every shortest path P between $w \notin T$ and $w' \in \mathrm{reg}(w,T) \cup T$, the simplified path S of P on $Q_1 \cup Q_2 \cup Q_3$ has the property that either*

(1) $S \subseteq \mathrm{reg}(w,T) \cup T$, or

(2) $S \cap T$ contains two nodes u,u' such that $S[w,u] \subseteq \mathrm{reg}(w,T) \cup T$, $S[u,u'] \subseteq \overline{\mathrm{reg}}(w,T) \cup T$, and $S[u',w'] \subseteq \mathrm{reg}(w,T) \cup T$.

Moreover, if $w' = x_i$, then $u \in Q_{(i \bmod 3)+1}$ and $S[u,u']$ induces at most two regions of $\overline{\mathrm{reg}}(w,T)$.

Proof. Let P be a shortest path between w and w' simplified on $T = Q_1 \cup Q_2 \cup Q_3$. If $P \subseteq \mathrm{reg}(w,T) \cup T$, we are done with P. Assume that P intersects $\overline{\mathrm{reg}}(w,T)$. Following the path P from w, let $z \in \mathbb{R}^2$ be the point closest to w that belongs to $P \cap \overline{\mathrm{reg}}(w,T)$, and let u be the node of $P \cap T$ just before z from w. By construction, $P[w,u] \subseteq \mathrm{reg}(w,T) \cup T$ is a shortest path. Similarly, define $z' \in \mathbb{R}^2$ as the point closest to w' that belongs to $P \cap \overline{\mathrm{reg}}(w,T)$, and u' as the node of $P \cap T$ just before z' from w'. By construction, $P[w',u'] \subseteq \mathrm{reg}(w,T) \cup T$ is a shortest path. Since $P[u,u']$ is simplified on T, we can check that $P[u,u']$ does not intersect any point of $\mathrm{reg}(w,T)$, hence $P[u,u'] \subseteq \overline{\mathrm{reg}}(w,T) \cup T$. Now, assume that $w' = x_i$. If $u \in (Q_{(i+1 \bmod 3)+1} \cup Q_i)$, then[1] P would be wholly contained in $\mathrm{reg}(w,T) \cup T$: a contradiction. Thus $u \in Q_{(i \bmod 3)+1}$. Moreover, if $P[u,u']$ induces more than two regions of $\overline{\mathrm{reg}}(w,T)$, then $P[u,u']$ is not simplified on T. ∎

Proposition 2.7 *Let T be a triangle defined by the node sequence (x_1, x_2, x_3) and the shortest paths (Q_1, Q_2, Q_3). Then, for every $w \notin T$, there exists a shortest path between w and either x_1 or x_2, which is entirely contained in* $\mathrm{reg}(w,T) \cup T$.

Proof. Assume, by way of contradiction, that no shortest path from w to either x_1 or x_2 is wholly contained in $\mathrm{reg}(w,T) \cup T$. For $i \in \{1,2\}$, let $P_i = P_i[w,u_i] \cup P_i[u_i,x_i]$ be a shortest path between w and x_i as described in Property 2.6, i.e., such that $P_i[w,u_i] \subseteq \mathrm{reg}(w,T) \cup T$ and $P_i[u_i,x_i] \subseteq \overline{\mathrm{reg}}(w,T) \cup T$. By assumption, P_1 and P_2 are not wholly contained in $\mathrm{reg}(w,T) \cup T$. The paths $P_1[u_1,x_1]$ and $P_2[u_2,x_2]$ must intersect at a point $z \in \overline{\mathrm{reg}}(w,T) \cup T$. Thus the shortest paths $P_1[w,z]$ and $P_2[w,z]$ have the same length, and so do the shortest paths $P_1[u_1,z] \cup P_2[z,x_1]$ and $Q_2[u_1,x_1]$ as well. Thus $P_1[w,u_1] \cup Q_2[u_1,x_1]$ is a shortest path wholly contained in $\mathrm{reg}(w,T) \cup T$; contradiction. ∎

The construction of the path B for a 3-gon H follows from Proposition 2.7. So, let us prove the second claim for $p = 4$. Let X be a shortest path between x_1 and x_3. Let $u \in X \cap (Q_1 \cup Q_4)$ be the farthest node from x_1. Let $z \in \mathbb{R}^2$ be the point of X just after u. Note that $z \notin H$. Let $u' \in X \cap (Q_1 \cup Q_3)$ be the farthest node from x_3. Note that the path $X[u',u'] \subseteq \mathrm{reg}(z,H)$. The path $Q_0 = X[x_1,u] \cup X[u,u'] \cup X[u',x_3]$ is a shortest path between x_1 and x_3 that is wholly contained in $\mathrm{reg}(z,H) \cup H$.

Consider the triangle T defined by the points (x_1, x_3, x_2) and the shortest paths (Q_0, Q_1, Q_2). By Proposition 2.7, there exists a shortest path P_1 between w and x_1 or x_3 wholly contained in $\mathrm{reg}(w,T) \cup T$. Without loss of generality assume that this path is between w and x_1. (The other case is symmetric, exchanging the roles of x_1 and x_3.) However, we have two cases to consider.

[1] Note that $(i + 1 \bmod 3) + 1$ is the predecessor of i in the cycle $(1, 2, 3)$.

Case 1: Q_0 is in $\mathrm{reg}(w, H) \cup H$, i.e., $\mathrm{reg}(z, H) = \mathrm{reg}(w, H)$.
Observe that $\mathrm{reg}(w, T) \subset \mathrm{reg}(w, H)$. Thus P_1 is also wholly contained in $\mathrm{reg}(w, H) \cup H$. We complete the proof setting $B = P_1$.

Case 2: Q_0 is in $\overline{\mathrm{reg}}(w, H) \cup H$, i.e., $\mathrm{reg}(z, H) = \overline{\mathrm{reg}}(w, H)$.
Consider the triangle T' defined by (x_1, x_3, x_4). Without loss of generality assume that P_1 is simplified on T'. (Indeed, since $P_1 \subseteq \mathrm{reg}(w, T) \cup T$, the simplified path of P_1 on T' is also wholly in $\mathrm{reg}(w, T) \cup T$.) If P_1 is wholly contained in $\mathrm{reg}(w, H) \cup H$, then we are done setting $B = P_1$. By Property 2.6, there exists $a \in P_1$ such that $P_1[w, a] \subseteq \mathrm{reg}(w, T') \cup T'$, and $P_1[a, x_1] \subseteq \overline{\mathrm{reg}}(w, T') \cup T'$. Moreover $a \in Q_3$. Without loss of generality assume the path is not wholly contained in $\mathrm{reg}(w, H) \cup H$ (so it intersects T'), and is simplified with T'.

Now consider the triangle T'' defined by (x_2, x_3, x_1). T'' is the same graph as T, except that its extremity has been permuted. By Proposition 2.7 there exists a shortest path between w and x_2 of x_3 wholly contained in $\mathrm{reg}(w, T'') \cup T''$. Without loss of generality we assume that the path is not wholly contained in $\mathrm{reg}(w, H) \cup H$, and is simplified with T'.

Case 2.1: The path is between w and x_3, denoted by P_3.
Because P_3 is simplified with T' and not wholly contained in $\mathrm{reg}(w, T') \cup T'$, by Property 2.6 there exists $u \in P_3 \cap Q_4$ such that $P_3[w, u] \subseteq \mathrm{reg}(w, T') \cup T'$ and $P_3[u, x_3] \subseteq \overline{\mathrm{reg}}(w, T') \cup T'$. Considering the closed curve $C = P_3[w, u] \cup Q_4[u, x_1] \cup P_1$, we check that $P_3[u, x_3]$ intersects P_1 at some node α. Depending on whether $\alpha \in P_1[w, a]$ or $\alpha \in P_1[a, x_1]$, the path $P_1[w, \alpha] \cup P_3[\alpha, x_3]$ or the path $P_1[w, a] \cup Q_3[a, x_3]$ is a shortest path between w and x_3 wholly contained in $\mathrm{reg}(w, H) \cup H$.

Case 2.2: The path is between w and x_2, denoted by P_2.
Let u, u' be the two nodes defined by Property 2.6 when applied on P_2 between w and x_2 w.r.t. the triangle T'. $u \in Q_3 \cup Q_4$. First, let us assume that $u \in Q_4$. Considering the closed curve defined by $P_2[w, u] \cup Q_4[u, x_1] \cup P_1$, we check that $P_2[u, x_2]$ intersects P_1 at some node α. If $\alpha \in P_1[w, a]$, then the path $P_1[w, \alpha] \cup P_2[\alpha, x_2]$ is a shortest path between w and x_2 wholly contained in $\mathrm{reg}(w, H) \cup H$. If $\alpha \in P_1[a, x_1]$, then either $u' \in Q_3$ or $u' \in Q_0$. If $u' \in Q_3$ the path $P_1[w, a] \cup Q_3[a, u'] \cup P_2[u', x_2]$ is a shortest path wholly contained in $\mathrm{reg}(w, H) \cup H$, and if $u' \in Q_0$ the path $P_1[w, a] \cup Q_3[a, x_3] \cup P_2[x_3, x_2]$ is a shortest path wholly contained in $\mathrm{reg}(w, H) \cup H$.

It remains to examine the case of $u \in Q_3$. Note that $u' \in Q_4$ or $u' \in Q_0$.

Because P_2 does not intersect $\overline{\mathrm{reg}}(w, T'')$ and P_2 is simplified on Q_0, we can assume without loss of generality that, if $u' \in Q_0$, then u' was chosen as the farthest node of $P_3 \cap Q_0$ from w (so $u' = x_1$). Doing by this way we force $u' \in Q_4$. We have to consider the shortest path P_3 between w and x_3, that is also assumed to be not wholly contained in $\mathrm{reg}(w, H) \cup H$, and simplified on T'. The segment of P_3 in $\mathrm{reg}(w, H) \cup H$ must intersect Q_4 (Q_2 and Q_3 are impossible). If P_3 intersects P_2 on the part $P_2[w, u]$ or $P_2[u', x_2]$ then we can easily construct with P_2 and P_3 a shortest path between w and x_2 or between w and x_3 that is wholly contained in $\mathrm{reg}(w, H) \cup H$. So, P_3 intersects $Q_4[u', x_4]$, say in a node b. Now, $P_3[b, x_3]$ (that is simplified with T') must intersect $P_2[w, u]$ or $P_2[u, u']$. In both cases we can construct a shortest path between w and x_3 that is wholly contained in $\mathrm{reg}(w, H) \cup H$, completing the proof of Case 2.2.

Therefore, in all the cases, we have constructed a shortest path B between w and some x_i, $i < p$, completing the proof of Lemma 2.5. ∎

Before proving the main result of the paper, we will also need the following easy lemma.

Lemma 2.8 *Let Q be a path of q nodes. For every integer $\lambda \geqslant 0$, $\gamma_\lambda(Q) = \lceil q/(2\lambda + 1) \rceil$.*

Proof. Let A be a minimum cardinality λ-dominating set of Q. Since each node of A dominates at most $2\lambda + 1$ nodes, we have $\gamma_\lambda(Q) \geqslant q/(2\lambda + 1)$. On the other hand; it is possible to split the q nodes of Q into $\lceil q/(2\lambda + 1) \rceil$ segments of $2\lambda + 1$ consecutive nodes (with the last segment possibly having fewer nodes). ∎

2.4 Main Result

In this section we prove our main result, namely, that $\gamma_k(G) < 3/\epsilon + 6$ for every planar graph G and $k \geqslant (5/7 + \epsilon)D$ (for any $\epsilon > 0$).

Towards proving this result, we introduce and study the related notion of a *k-dominating subgraph*. Given a graph G, a connected subgraph H is called a k-dominating subgraph of G if $V(H)$ is a k-dominating set for G.

Our main result is derived from the following technical theorem concerning k-dominating subgraphs, which may be of independent interest.

Theorem 2.9 *For every planar graph G of diameter D, and for every $k \geqslant \lfloor 5D/7 \rfloor$, there exists a k-dominating subgraph H of G such that H is composed of at most 6 shortest paths of G, and $|V(H)| \leqslant 6D - 1$.*

Proof. We present a constructive proof, implying a polynomial time algorithm for the construction of such a k-dominating subgraph H. Without loss of generality we assume hereafter that $k < D$, as the result clearly holds for $k = D$. The algorithm constructs H gradually. Letting $\delta(H)$ denote the *deviation* of a candidate subgraph H from being a k-dominating subgraph, i.e., $\delta(H) = |V \setminus \Gamma_k(H)|$, we attempt to progress by repeatedly decreasing $\delta(H)$.

Step 1: We start assuming that $k \geqslant \lfloor D/2 \rfloor$. Let x_1 and x_2 be two nodes such that $d_G(x_1, x_2) > k$, and let Q_1 be any shortest path connecting them.

If $V(Q_1)$ is k-dominating, then taking $H = Q_1$ we are done, as H is connected and $|V(H)| \leqslant D$.

Step 2: We construct a 3-gon H for G as follows. Let $x_3 \notin \Gamma_k(Q_1)$, and let Q_2 (resp., Q_3) be a shortest path between x_3 and x_1 (resp., between x_3 and x_2), chosen such that $H = Q_1 \cup Q_2 \cup Q_3$ induces at most two regions of $\mathbb{R}^2$. This can be done by ensuring that for every $t \in Q_2 \cap Q_3$, the subpaths of Q_2 and Q_3 from x_3 to t coincide, i.e., $Q_2[x_3, t] = Q_3[x_3, t]$, and analogously, that for every $t \in Q_2 \cap Q_1$, $Q_2[x_1, t] = Q_1[x_1, t]$, and that for every $t \in Q_3 \cap Q_1$, $Q_3[x_2, t] = Q_1[x_2, t]$. One can readily verify that H is a 3-gon of G.

Step 3: H is a 3-gon.
If $V(H)$ is a k-dominating set then we are done, as H is connected and $|V(H)| \leqslant 3D$.

Otherwise, let $w \notin \Gamma_k(H)$. By Lemma 2.5, G has a $(3 - i + 2)$-gon and a $(i + 2)$-gon. Since $1 \leqslant i < 3$, G has a 4-gon. Set H to be this 4-gon.

Step 4: H is a 4-gon.
If $V(H)$ is a k-dominating set, then we are done, as H is connected and $|V(H)| \leqslant 4D$.

Otherwise, let $m = \delta(H) > 0$, and let $w \notin \Gamma_k(H)$. Assume that H is defined by the nodes $(x_1, \ldots, x_4)$ and the shortest paths $(Q_1, \ldots, Q_4)$. By

Lemma 2.5, let A be the shortest path between w and x_4 that is wholly contained in $\mathrm{reg}(w,H) \cup H$, and let B be the second shortest path between w and x_i, $i < 4$, that is wholly contained in $\mathrm{reg}(w,H) \cup H$. Let $H_1 = B \cup (Q_i \cup Q_{i+1} \cup \cdots \cup Q_3) \cup A$ and $H_2 = B \cup (Q_{i-1} \cup \cdots \cup Q_1) \cup Q_4 \cup A$ be the respective $(6-i)$-gon and $(i+2)$-gon whose existence is asserted by Lemma 2.5. Note that H_1 and H_2 are 3-, 4- or 5-gons. Figure 3 depicts H_1 when $i = 2$, in which case $H_1 = B \cup Q_2 \cup Q_3 \cup A$, namely, it is obtained by following the cycle w, x_2, x_3, x_4, w.

Let $H' = H_1 \cup H_2 = A \cup H \cup B$. If $V(H')$ is a k-dominating set, then we are done, as H' is connected and $|V(H')| \leqslant 6D-1$. Indeed, in the worst-case H' is composed of 6 paths of length D that meet in 5 nodes ($x_1, \ldots, x_4$ and w), thus has $5 + 6(D-1) = 6D-1$ nodes. So hereafter assume that $V(H')$ is not k-dominating. Let $m' = \delta(H') > 0$, and let $w' \notin \Gamma_k(H')$. Note that $m' \leqslant m-1$ because $V(H')$ covers $\Gamma_k(H)$ and also w (as $V(H) \subset V(H')$).

H' defines at most three regions, namely, $\overline{\mathrm{reg}}(w,H)$, and $\mathrm{reg}(w,H)$, which in turn is split by H_1 into two sub-regions, $R_1 = \overline{\mathrm{reg}}(z,H_1)$ and $R_2 = \mathrm{reg}(w,H) \setminus (R_1 \cup H_1)$, where z is an arbitrary node of $V(H_1) \setminus V(H_2)$ (such a node must exist). See Fig. 3. Note that on Fig. 3 two incident paths may partially overlap: for instance $A[w,t] = B[w,t]$ for some t is possible. The proof anyway is not based on this particular example.

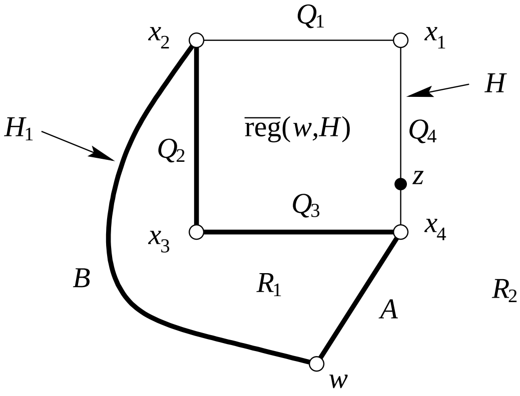

Figure 3: The ψ-gon H_1 (in bold edges). A node $w' \notin \Gamma_k(H')$ has to belong to R_1 or R_2.

Corollary 2.2 implies that $w' \notin \overline{\mathrm{reg}}(w,H)$, since w and w' must belong to the same region. So assume that $w' \in R_j$, for $j = 1$ or 2. Note that

$R_j = \mathrm{reg}(w', H_j)$. We now split the discussion into two cases.

Case 1: H_j is a 3-gon or a 4-gon.
In this case, we prove that $\Gamma_k(H_j) = \Gamma_k(H')$. Clearly, $\Gamma_k(H_j) \subseteq \Gamma_k(H')$. The converse is proved by contradiction. By Corollary 2.2, H_j covers all the points of $\overline{\mathrm{reg}}(w', H_j)$, and thus H' covers those points as well since $H_j \subset H'$. Assume, to the contrary, that there is some $w'' \in R_j$ such that $w'' \in \Gamma_k(H')$ but $w'' \notin \Gamma_k(H_j)$. This implies that there exists some $u \in V(H') \setminus V(H_j)$ such that $w'' \in \Gamma_k(u)$. Any shortest path from u to w'' must intersect H_j at some node v because H_j is a separator. Thus $w'' \in \Gamma_k(v)$, proving that $w'' \in \Gamma_k(H_j)$, contradiction. Thus $\Gamma_k(H_j) \supseteq \Gamma_k(H')$ as well.

It follows that the number of nodes uncovered by H_j is $m' \leqslant m - 1$. Depending on whether H_j is a 3-gon or a 4-gon, we now set $H \leftarrow H_j$ and go to Step 3 or Step 4 and reiterate, and by induction the number of uncovered nodes $\delta(H)$ will decrease to 0.

Case 2: H_j is a 5-gon.
Note that in this case the shortest path B is between w and x_1 or x_3, w.l.o.g. x_1. (The other case is handled symmetrically by exchanging node names x_1 and x_3.)

Assume there is a path B' between w and x_2 that is wholly contained in $\mathrm{reg}(w, H) \cup H$. Then, we can remove the path B constructed while applying Lemma 2.5, and rename B' into B. The ψ-gons H_1 and H_2 resulting from considering the new simplified paths A and B are necessarily 4-gons, and we can conclude by Case 1.

Similarly, if there is a path A' between w and x_3 that is wholly contained in $\mathrm{reg}(w, H) \cup H$, then we can remove path A and rename A' into A. The 4-gons resulting from considering the new simplified paths A and B allow us to go again to Case 1.

We are thus left with the situation when none of the shortest paths from w to x_2 and from w to x_3 is wholly contained in $\mathrm{reg}(w, H) \cup H$.

Let X be a shortest path from w to x_2, and let Y be a shortest path from w to x_3. (Without loss of generality we assume that X and Y are simplified on H). By assumption, X and Y are not wholly contained in $\mathrm{reg}(w, H) \cup H$. Therefore, X must intersect Q_3 or Q_4, and Y must intersect Q_1 or Q_4. Let u be the farthest node from w such that $u \in X \cap (Q_3 \cup Q_4)$, and let v be the farthest node from w such that $v \in Y \cap (Q_1 \cup Q_4)$.

Case 2.1: $u \in Q_3$.
Similarly to the proof of Lemma 2.5, X and Y intersect at some point $\alpha \in \overline{\mathrm{reg}}(w, H)$. (See Fig. 4.) Then $X[w, u]$ is wholly contained in $\mathrm{reg}(w, H) \cup H$ and intersects Q_3 leading to x_3. As both X and Y induce shortest paths from w to α, $|X[w, u]| + |X[u, \alpha]| = |Y[w, \alpha]|$, or $|X[w, u]| = |Y[w, \alpha]| - |X[u, \alpha]|$. Also, $|Q_3[u, x_3]| \leqslant |X[u, \alpha]| + |Y[\alpha, x_3]|$. Hence $|X[w, u]| + |Q_3[u, x_3]| \leqslant |Y[w, \alpha]| + |Y[\alpha, x_3]| = |Y[w, x_3]|$, so $X[w, u] \cup Q_3[u, x_3]$ is a shortest path from w to x_3 wholly contained in $\mathrm{reg}(w, H) \cup H$, contradicting the definition of Y. Thus Case 2.1 is not possible.

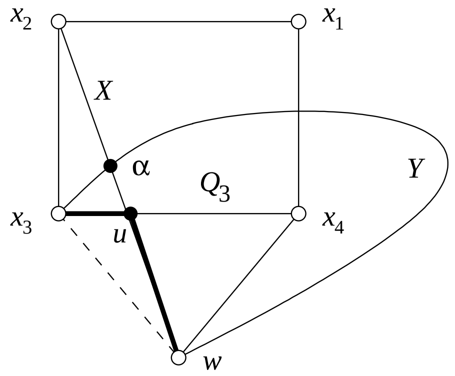

Figure 4: Case 2.1.

Case 2.2: $v \in Q_1$.
This case is symmetric to the previous one; the path $Y[w, v] \cup Q_1[v, x_2]$ is a shortest path wholly contained in $\mathrm{reg}(w, H) \cup H$, contradicting the definition of X.

Case 2.3: $u, v \in Q_4$.
In this case we are left with a configuration similar to the one depicted on Fig. 5.

We now assume that $k \geqslant \lfloor 5D/7 \rfloor$. Let $a = d_G(x_1, u)$, $b = d_G(v, x_4)$ and $c = d_G(u, v)$. Without loss of generality we assume that $a \leqslant b$. Note that since $a + c + b \leqslant D$, we have $a \leqslant (D - c)/2$. Note also that we have $d_G(u, x_2) \leqslant D - \lfloor k \rfloor - 1$ and $d_G(v, x_3) \leqslant D - \lfloor k \rfloor - 1$ because $w \notin \Gamma_k(u)$ and $w \notin \Gamma_k(v)$.

If $c \geqslant D/7$, then

$$d_G(x_1, x_2) \leqslant a + d_G(u, x_2) \leqslant \frac{1}{2}(D - c) + D - \lfloor k \rfloor - 1$$

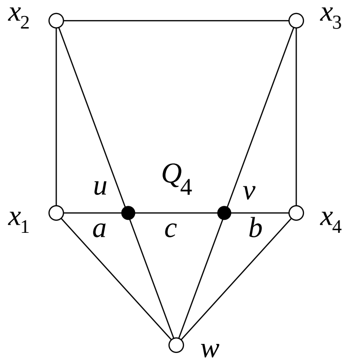

Figure 5: The final configuration.

$$\leqslant \ \frac{1}{2}\left(D - \frac{1}{7}D\right) + D - \frac{5}{7}D + 1 - 1 \ = \ \frac{5}{7}D - 1 \ . \qquad (1)$$

It follows that $d_G(x_1, x_2) \leqslant 5D/7 - 1 < k$: a contradiction since H is a 4-gon, forcing to have $d_G(x_1, x_2) > k$.

If $c \leqslant D/7$, then

$$\begin{aligned} d_G(x_2, x_3) \ &\leqslant \ d_G(x_2, u) + c + d_G(v, x_3) \ \leqslant \ 2D - 2\lfloor k \rfloor - 2 + \frac{1}{7}D \\ &\leqslant \ 2D - 2\left(\frac{5}{7}D\right) - 2 + \frac{1}{7}D \ = \ \frac{5}{7}D - 2 \ . \qquad (2) \end{aligned}$$

It follows that $d_G(x_2, x_3) \leqslant 5D/7 - 2 < k$, leading to contradiction since H is a 4-gon, forcing to have $d_G(x_2, x_3) > k$.

Hence the Case 2.3 is impossible as well. The theorem follows. ▮

We are now ready to state and prove our main result.

Theorem 2.10 *For every planar graph G of diameter D, for every $\epsilon > 0$, and for every integer $k \geqslant (5/7 + \epsilon)D$, we have $\gamma_k(G) < 3/\epsilon + 6$.*

Proof. The connected subgraph H constructed in Theorem 2.9 consists of the union of at most 6 shortest paths, each of length at most D, and is a k-dominating set for G for $k \geqslant \lfloor 5D/7 \rfloor$. Choosing a minimum cardinality λ-dominating set for each path, we obtain a $(k+\lambda)$-dominating set B for G. By Lemma 2.8, $|B| \leqslant 6\lceil D/(2\lambda + 1) \rceil$.

Let us set $\lambda = \lceil \epsilon D \rceil$. B is therefore a k'-dominating set for G for every integer $k' \geqslant \lfloor 5D/7 \rfloor + \lceil \epsilon D \rceil \geqslant (5/7+\epsilon)D$. We have, for $\epsilon > 0$,

$$|B| \;\leqslant\; 6\left\lceil \frac{D}{2\lambda+1} \right\rceil \;<\; \frac{6D}{2\lceil \epsilon D \rceil + 1} + 6 \;<\; \frac{6D}{2\epsilon D+1} + 6 \;<\; \frac{3}{\epsilon} + 6,$$

completing the proof of Theorem 2.10. ∎

2.5 Conjecture and Worst-Case

We leave open the problem of bounding $\gamma_k(G)$ for $k \geqslant D/2$, and make the following conjecture.

Conjecture 2.11 *There exists a constant c_0 such that for every planar graph G of diameter D, and for every $k \geqslant D/2$, $\gamma_k(G) \leqslant c_0$.*

Here we show that $c_0 \geqslant 4$.

Theorem 2.12 *For every even diameter $D \geqslant 8$, there exists a planar graph G of diameter D for which $\gamma_{D/2}(G) = 4$.*

Proof. Let $D = 2t$ for $t \geqslant 4$. The graph denoted by G_t is composed of

(1) two nodes A and B called *poles*;

(2) D disjoint paths $P_1, \ldots, P_D$ called *meridians*, each of length D, joining A to B;

(3) a cycle of length $2D$, called *equator*, joining the middle nodes of each meridians and with exactly one node of degree two between two consecutive meridians.

(See Fig. 6 for $t = 4$).

Observe that for a cycle C_{2D} of length $2D$, $\gamma_{D/2}(C_{2D}) = 2$. Hence taking $S = \{A, B\}$, it is clear that all the nodes of the meridians are covered, but the degree 2 nodes of the equator are note covered. Thus to complete S into a dominating set, we have to add two opposite nodes u, v on the equator.

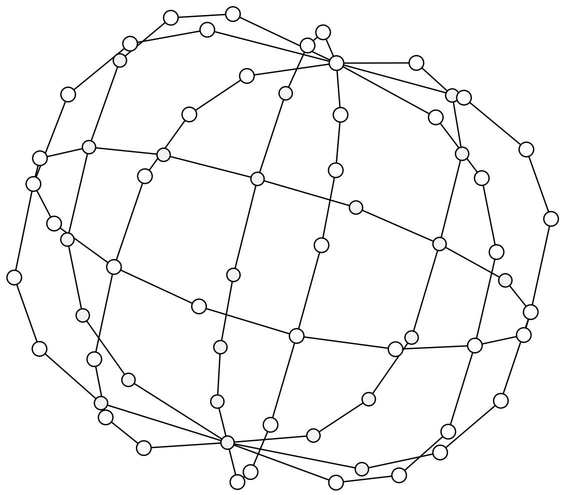

Figure 6: The planar graph embedded on a sphere for $D = 8$ such that $\gamma_4(G) = 4$.

The resulting set $S = \{A, B, u, v\}$ is then a $D/2$-dominating set of G_t. It follows that $\gamma_{D/2}(G_t) \leqslant 4$.

Now if we take two opposite nodes different from A and B belonging to two distinct meridians we will not cover all the nodes of the $2t$ meridians. In this case we will need more than 4 nodes ($t \geqslant 4$) to cover all the nodes of the meridians. This completes the proof of the lower bound. ∎

3 Treewidth and Chordality Bounded Graphs

In this section we establish some additional results concerning the value of $\gamma_k(G)$ for various graph classes. We use the notion of tree-decomposition introduced by Robertson and Seymour in their work on graphs minors [RS86]. A *tree-decomposition* of a graph $G = (V, E)$ is a tree $T = (\tilde{V}, \tilde{E})$ whose nodes are subsets of V, such that

1. $\bigcup_{X \in \tilde{V}} X = V$;
2. for every edge $(u, v) \in E$, there exists a node set $X \in \tilde{V}$ such that $u, v \in X$; and

3. for every $X, Y, Z \in \tilde{V}$, if Y is on the path from X to Z in T, then $X \cap Z \subseteq Y$.

The *width* of the tree-decomposition T is $\omega(T) = \max_{X \in \tilde{V}}\{|X| - 1\}$. The *treewidth* of a graph G is $\min_T\{\omega(T)\}$, where the minimum is taken over all the tree-decompositions T of G. Note that trees have treewidth 1. We use the following property whose proof can be founded in [Die00, pp. 256].

Property 3.1 *Let $\{X_1, X_2\}$ be any edge of a tree-decomposition T of G, and let T_1, T_2 be the components of $T \setminus \{X_1, X_2\}$, with $X_1 \in T_1$ and $X_2 \in T_2$. Then $X_1 \cap X_2$ separates $U_1 = \bigcup_{A \in T_1} A$ from $U_2 = \bigcup_{B \in T_2} B$ in G.*

Theorem 3.2 *For every graph $G = (V, E)$ of diameter D and treewidth t, and for every $k \geqslant \lfloor D/2 \rfloor$, we have $\gamma_k(G) \leqslant t + 1$.*

Proof. Let $T = (\tilde{V}, \tilde{E})$ be a width t tree-decomposition of G, i.e., such that $|X| \leqslant t + 1$ for every $X \in \tilde{V}$. We prove the theorem by showing that there exists some set $X^* \in \tilde{V}$ that k-dominates G.

For $X \in \tilde{V}$, let $T(X)$ denote the subtree of T rooted at X, and let $S(X) = \bigcup_{Y \in T(X)} Y$. Also let $\mathrm{Far}(X) = V \setminus \Gamma_k(X)$.

We construct a cycle-free path $X_1, X_2, \ldots, X_m$ in T step by step, such that X_1 is the root of T, $X_m = X^*$, and X_i does not k-dominate G for $i < m$. The initial step consists of selecting X_1 and checking whether it k-dominates G or not. Starting with a current node X_i of T, Step i consists of choosing the unique child X_{i+1} of X_i such that $S(X_{i+1})$ contains $\mathrm{Far}(X_{i+1})$. It suffices to prove by induction on i that Step i is doable, namely, that there exists such a unique child X_{i+1}. Indeed, since the length of a cycle-free path in T is finite, we stop the process, in the worst-case, once reaching a leaf of T.

Let $U = \mathrm{Far}(X_i)$, and assume $U \neq \varnothing$. Observe that X_i is a separator of G by Property 3.1. Let us show that there is one connected component of $T \setminus \{X_i\}$ that contains all the nodes of U. Assume the opposite, and let A, B be two nodes of two different connected components of $T \setminus \{X_i\}$ containing each a node of U, say, $a \in A \cap U$ and $b \in B \cap U$. If $a = b$, then $a \in X_i$ by the definition of the tree-decomposition, contradicting the fact that $a \in U$ and $U \cap X_i = \varnothing$. On the other hand, if $a \neq b$, then by Proposition 3.1 a and b

are in different components of G. In the other hand, by Proposition 2.1 any shortest path from a to b in G cannot intersect X_i, contradicting the fact that X_i is a separator of G.

So, let $\{X_i, X_{i+1}\}$ be the edge that leads to the unique connected component of $T \setminus \{X_i\}$ that contains all the node of U. Note that X_{i+1} is a child of X_i (by induction, all the nodes of U were in the subtree of root X_i). Let $I = X_i \cap X_{i+1}$, and let $U' = V(G) \setminus \Gamma_k(I)$. The key observation is that U' is also contained in the subtree rooted at X_{i+1}. Indeed, consider any node A of the subtree of root X_{i+1} and any node B from another connected component of $T \setminus \{X_{i+1}\}$, and let $a \in A \cap U'$ and $b \in B \cap U'$. If $a = b$, then by the definition of tree-decomposition $a \in X_i$ and $a \in X_{i+1}$, hence $a \in I$, contradicting the fact that by choice $a \notin I$. If $a \neq b$, then observe that I separates a and b in G (by Proposition 3.1), leading to contradiction with Proposition 2.1, by which any shortest path from a to b cannot intersect I.

Obviously, completing I by some extra nodes of X_{i+1}, in order to obtain X_{i+1}, we keep the induction property, and prove by this way that Step i is doable. This completes the proof. ∎

Theorem 3.2 is the best we can hope for, in the sense that for every t, and every odd $D \geqslant 3$, there exists a graph G of diameter D and treewidth t such that $\gamma_k(G) = t + 1$, for $k \geqslant \lfloor D/2 \rfloor$. Consider for instance a clique K_{t+1} modified by attaching a distinct path p_i of length ℓ, for $\ell \geqslant 1$, to each node i in the clique. The diameter of the resulting graph is $D = 2\ell + 1$, its treewidth is t, and $\gamma_\ell(G) = t + 1$, as any ℓ-dominating set must contain at least one distinct node for dominating each of the $t + 1$ paths p_i.

It is known that graphs of genus g and diameter D have a treewidth bounded by $O(gD)$, cf. [Epp00]. For planar graph one can tighten this bound. For that we need the following definition. A plane graph is a 1-*outerplanar embedding* if all its nodes lie on the outer face. For $r \geqslant 2$, a plane graph is an r-*outerplanar embedding* if deleting all the nodes of the outer face yields an $(r-1)$-outerplanar embedding of the remaining graph. A graph is r-outerplanar if it has a r-outerplanar embedding. We refer to 1-outerplanar graphs (such as trees and $K_{2,3}, K_4$-free graphs) simply as outerplanar graphs.

Theorem 3.3 *For every planar G of diameter D, and for every $k \geqslant \lfloor D/2 \rfloor$, we have $\gamma_k(G) \leqslant 3D + 3$.*

Proof. Assume that G is r-outerplanar for $r \geqslant 1$. We first assume that $r \geqslant 2$. Consider an r-outerplanar embedding of G. Let V_1 denote the outer layer of G, namely, the set of all the nodes lying on the outer face of G. Then, successively for $i = 2, \ldots, r$, let V_i denote the set of all the nodes lying on the outer face of $G \setminus (V_1 \cup \cdots \cup V_{i-1})$. For every i, consider some $u_i \in V_i$. Observe that for every $j > i$, we have $d_G(u_i, u_j) \geqslant j - i$. Indeed, if $d_G(u_i, u_j) < j - i$, then there is a path of at most $j-i-1$ edges linking a node of V_i to a node of V_j, which is impossible for a plane embedding of G (there is no edge between a node of V_i and a node of V_{i+2}). Thus $d_G(u_1, u_r) \geqslant r - 1$, implying that $D \geqslant r - 1$. Obviously, the latter result holds for $r = 1$.

Now, in [Bod98], it is shown that r-outerplanar graphs have treewidth at most $3r - 1$. The claim follows by Theorem 3.2. ∎

A graph is *c-chordal* if every induced cycle is of length at most c. Triangulated graphs are precisely 3-chordal graphs. Note that a c-chordal graph does not necessarily have a small separator; in fact, the size of its smallest separator is at least the size of its maximum clique. So, Theorem 3.2 might be not very useful. However, we can show the following.

Theorem 3.4 *For every c-chordal graph G of diameter D, and for every $k \geqslant \lfloor D/2 \rfloor + \lfloor c/2 \rfloor$, we have $\gamma_k(G) = 1$.*

Proof. It is proved in [GKK$^+$00] that c-chordal graphs have a tree-decomposition $T = (\tilde{V}, \tilde{E})$ (computable in polynomial time) such that for every $X \in \tilde{V}$ and all $u, v \in X$, $d_G(u, v) \leqslant c/2$. Let $k_1 = \lfloor D/2 \rfloor$ and $k_2 = \lfloor c/2 \rfloor$. As shown in the proof of Theorem 3.2, there exists some $X^* \in \tilde{V}$ that k_1-dominates G. Let $u \in X^*$. As $d_G(u, v) \leqslant k_2$ for every $v \in X^*$, necessarily $\Gamma_{k_2}(u) \supseteq X^*$. It follows that $\{u\}$ is a $(k_1 + k_2)$-dominating set of G. ∎

4 Application to Compact Routing

Let us now switch to discussing an application of our results in the area of communication networks. A point-to-point communication network is modeled as a graph $G = (V, E)$, where the set of nodes represent the processors of the network and every pair of two opposite arcs represents a bidirectional communication link. A *routing scheme* R is a distributed algorithm whose role is to deliver messages between nodes of the network. The routing scheme consists of certain *distributed data structures* in the network, and a *delivery protocol*, which can be invoked in any node u with two parameters: a *routing label* $\mathcal{L}(v)$ of the destination node v, and the message's information field. The message is delivered to v via a sequence of transmissions determined uniquely by the distributed data structure, i.e., at each intermediate node along the route, the routing scheme decides the next neighbor to which the message should be forwarded. The length of the route traversed by a message from u to v in the graph G according to the routing scheme R is denoted by $\rho_R(u, v)$. The *dilation* of a routing scheme R is the maximal route length of a path traversed by a message, formally $\max_{u \neq v} \{\rho_R(u, v)\}$.

An *interval routing scheme* R on G is a routing scheme consisting of a pair $(\mathcal{L}, \mathcal{I})$, generated in the preprocessing step, where $\mathcal{L}$ is a *node-labeling*, $\mathcal{L} : V \to \{1, \ldots, n\}$, and $\mathcal{I}$ is an *arc-labeling*, $\mathcal{I} : E \to 2^{\mathcal{L}(V)}$, that satisfy the following condition. For any node u, the collection of sets that label all the outgoing arcs of u forms a partition of the name range (possibly excluding u itself[2]). Formally, for every $u \in V$, denoting by E_u the set of arcs incident to u,

1. $\bigcup_{e \in E_u} \mathcal{I}(e) = \{1, \ldots, n\} \setminus \{\mathcal{L}(u)\}$,
2. $\mathcal{I}(e_1) \cap \mathcal{I}(e_2) = \varnothing$ for every two distinct arcs $e_1, e_2 \in E_u$.

The delivery protocol is defined as follows. Given a destination node v, set the first header to be $h = \mathcal{L}(v)$. Also, for every node u, receiving a message with header h, first check if $\mathcal{L}(u) = h$, and deliver the message locally if equality holds (meaning that u is the message's final destination). If not, then forward the message with the same header h on the output port corresponding to the unique arc (u, w) such that $h \in \mathcal{I}(u, v)$, namely, the arc whose labeling set contains the destination label.

Note that any interval routing scheme can be implemented by classical routing tables. The main difference is in the coding of the table.

[2]A labeling excluding u from its arc-labels is termed *strict*. Although non-strict labeling may produce more compact schemes, in this paper we restrict our attention to strict labeling only.

Given an integer n and a subset $I \subseteq \{1, \ldots, n\}$, define the *compactness of I w.r.t. n*, denoted $c_n(I)$, as the smallest integer k such that I can be represented by the union of k intervals $[a, b]$ of consecutive integers from $\{1, \ldots, n\}$, with n and 1 being considered as consecutive (cyclically). The *compactness* of an interval routing scheme $R = (\mathcal{L}, \mathcal{I})$ on G is the maximum, over all arcs $e \in E$, of the compactness $c_n(\mathcal{I}(e))$ of the set $\mathcal{I}(e)$ labeling e. The *total compactness* is the sum $\sum_{e \in E} c_n(\mathcal{I}(e))$.

Because interval routing schemes can be implemented by routing tables, we deal with the compactness of a routing table. Intuitively, smaller compactness and degrees imply smaller routing tables. The interval routing strategy presented in [SK85], based on routing on a minimum spanning tree, has compactness 1 (for every graph of diameter D) and dilation at most $2D$. It is known that there are worst-case graphs for which every routing table of compactness 1 has dilation at least $2D - 3$ [TL97]. It is also known that every graph has routing table of compactness $\sqrt{n \ln n} + O(1)$ with dilation at most $\lceil 1.5D \rceil$ [KRŠ00], whereas there are worst-case graphs for which every routing table of compactness k has dilation at least $\lfloor 1.5D \rfloor - 1$ for every $k = \Omega(\frac{n}{D \log(n/D)})$ [Gav00a]. Other results can be found on the recent survey [Gav00b].

Here we show a trade-off between dilation and compactness.

Theorem 4.1 *Let G be a graph of diameter D with n nodes and m edges. Let $k \geqslant 0$, and let $t = \gamma_k(G)$. Then, G has a routing table with dilation at most $D + k$, compactness at most $t/2 + 1$, and total compactness at most $tn + 2m$. Moreover, if $t/\log n \to +\infty$, then the compactness can be reduced to $t/4 + o(t)$.*

Proof. Let $\{x_1, \ldots, x_t\}$ be a k-dominating set in G of minimal cardinality $t = \gamma_k(G)$. Construct a partition of the graph nodes into t pairwise disjoint connected regions $R_i \subseteq \Gamma_k(x_i)$ around each x_i, such that each region R_i consists of nodes closest to x_i (breaking tie arbitrarily) and the union of the regions covers all the nodes. Obviously the radius of each region is at most k.

For every x_i, construct a shortest path spanning tree T_i for G. Let $\hat{T}_i$ be the restriction of T_i to the region R_i, i.e., $\hat{T}_i = T_i \cap R_i$.

Partition the range of integers $[1, n]$ into t contiguous intervals I_i of size $|R_i|$, for $i \in \{1, \ldots, t\}$. For each region R_i, assign each node $v \in R_i$ a distinct label $\mathcal{L}(v)$ from the range I_i, in DFS order, starting at x_i. Use these labels to define an interval routing scheme on $\hat{T}_i$ as in [SK85]. The compactness is 2 (instead of 1) because a cyclic interval in I_i has to be represented by two sub-intervals of I_i. For every two nodes $u, w \in R_i$, this scheme yields a shortest route on $\hat{T}_i$ (albeit perhaps not shortest in G).

In addition, for every $1 \leqslant i \leqslant t$ and for every node $w \notin R_i$, add the interval I_i to the arc connecting w to its parent in T_i.

As the regions are disjoint, it is easy to verify that for every node v, all the intervals assigned to the arcs of v are pairwise disjoint. Overall, at most $t/2+1$ intervals were assigned to each arc because every subset of at most t intervals has no more than $t/2$ non contiguous intervals (one more sub-interval is required for the own region of v). In fact, exactly $t+\deg(v)$ intervals altogether ($\deg(v)+1$ intervals for R_i and $t-1$ for all the others) are assigned to the arcs of a node v of degree $\deg(v)$. Thus the total compactness is $tn+2m$.

Observe that the resulting route between nodes belonging to the same region has length at most $2k \leqslant D+k$. As for routing from a node $u \in R_i$ to a node $w \in R_j$, $j \neq i$, note that the first segment of the route, until reaching R_j, proceeds along a shortest path from u to x_j, and once entering R_j, the remaining segment of the route follows a shortest path. Hence the total length of the route cannot exceed $d_G(u,x_j)+d_G(x_j,w) \leqslant D+k$.

Actually, the bound on the compactness can be slightly reduced to $t/4+o(t)$, for t large enough, using the result of [GP99]. This result shows that, for every t, there exists a permutation π of $\{1,\ldots,t\}$ such that the compactness of every set F of a family $\mathcal{F}$ of subsets of $\{1,\ldots,t\}$, with $|\mathcal{F}| < \exp(t/2)/t$, satisfies:

$$c_t(\pi(F)) \;<\; \frac{t}{4}+\frac{1}{4}\sqrt{2t\ln(|\mathcal{F}|\cdot n)},$$

where $\pi(F)$ denotes the set $\{\pi(x) \mid x \in F\}$. Let $\mathcal{I}(e)$ be the set of labels assigned to arc e by the previous interval routing scheme. Set $\mathcal{F} = \{F(e) \mid e \in E(G)\}$, where $F(e) = \{i \mid I_i \subseteq \mathcal{I}(e)\}$ corresponds to the set of region's indices contained in $\mathcal{I}(e)$. Note that the range of labels of each region (except for v's) is wholly contained in $\mathcal{I}(e)$. Since $\bigcup_{j=1}^{t} I_j = [1,n]$, we have that $c_n(\mathcal{I}(e)) \leqslant c_t(F(e))+1$. Observe that the relative order of the ranges can be chosen arbitrary: all the labels of some I_j can be chosen to be larger or smaller than all the labels of any other range. So, indices of I_j's ranges can be permuted by some permutation π of $[1,t]$ in order to decrease $c_t(\pi(F(e)))$. By [GP99], if $|\mathcal{F}| < \exp(o(t))/t$, then there exists π_0 such that $c_t(\pi_0(F(e))) < t/4+o(t)$, for every $F(e) \in \mathcal{F}$. Note that $|\mathcal{F}| = 2m$. So, the condition $|\mathcal{F}| < \exp(o(t))/t$ can be replaced by $t/\log n \to +\infty$, as $m = O(n^2)$. ∎

This result can be viewed as a generalization of [KRŠ00], originally proved for $k = \lceil D/2 \rceil$. In that paper it was shown that every graph G satisfies $\gamma_k(G) \leqslant \lceil \sqrt{n \ln n}\, \rceil$, and the compactness derived therein was $\lceil \sqrt{n \ln n}\, \rceil + 1$. Thus Theorem 4.1 improves the compactness bound obtained in [KRŠ00] by an asymptotic factor of 4. The same kind of construction has been also implicitly used in [TL99]. Specifically, it is shown in [TL99] that every graph G satisfies $\gamma_k(G) \leqslant \sqrt{n}$ for $k \geqslant 2D/3$. Thus, every graph has routing table of compactness at most $\sqrt{n}/4 + o(\sqrt{n})$ for dilation $\lceil 5D/3 \rceil$. (The original compactness result of [TL99] was $\lceil \sqrt{n}\, \rceil + 1$.)

Combining Theorem 2.10 and Theorem 4.1, we obtain the following.

Corollary 4.2 *Let $\epsilon > 0$ be an arbitrary constant. Every planar graph has a routing table with dilation at most $\lceil (12/7 + \epsilon)D \rceil$, and constant compactness.*

Note that Conjecture 2.11 for $c_0 = 4$ implies that every planar graph of diameter D enjoys a routing table with dilation at most $\lceil 1.5D \rceil$ and compactness at most 3.

References

[Bod98] Hans Leo Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical Computer Science*, 209:1–45, 1998.

[CV00] Victor Chepoi and Yann Vaxes. On covering bridged plane triangulations with balls. Manuscript submitted, 2000.

[Die00] Reinhard Diestel. *Graph Theory (second edition)*, volume 173 of Graduate Texts in Mathematics. Springer, February 2000.

[Epp00] David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27:275–291, 2000.

[Gav00a] Cyril Gavoille. On the dilation of interval routing. *The Computer Journal*, 43(1):1–7, 2000.

[Gav00b] Cyril Gavoille. A survey on interval routing. *Theoretical Computer Science*, 245(2):217–253, 2000.

[GKK+00] Cyril Gavoille, Michal Katz, Nir A. Katz, Christophe Paul, and David Peleg. Approximate distance labeling schemes. Research Report RR-1250-00, LaBRI, University of Bordeaux, 351, cours de la Libération, 33405 Talence Cedex, France, December 2000.

[GP99] Cyril Gavoille and David Peleg. The compactness of interval routing. *SIAM Journal on Discrete Mathematics*, 12(4):459–473, October 1999.

[KRŠ00] Rastislav Kráľovič, Peter Ružička, and Daniel Štefankovič. The complexity of shortest path and dilation bounded interval routing. *Theoretical Computer Science*, 234(1-2):85–107, 2000.

[Lov75] Laszlo Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.

[MS96] G. MacGillivray and K. Seyffarth. Domination numbers of planar graphs. *Journal of Graph Theory*, 22(3):213–229, 1996.

[RS86] Neil Robertson and Paul D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7:309–322, 1986.

[SK85] Nicola Santoro and Ramez Khatib. Labelling and implicit routing in networks. *The Computer Journal*, 28(1):5–8, February 1985.

[TL97] Savio S. H. Tse and Francis C. M. Lau. An optimal lower bound for interval routing in general networks. In 4^{th} *International Colloquium on Structural Information & Communication Complexity (SIROCCO)*, pages 112–124. Carleton Scientific, July 1997.

[TL99] Savio S. H. Tse and Francis C. M. Lau. Some results on the space requirement of interval routing. In 6^{th} *International Colloquium on Structural Information & Communication Complexity (SIROCCO)*, pages 264–279. Carleton Scientific, July 1999.

[vLT87] Jan van Leeuwen and Richard B. Tan. Interval routing. *The Computer Journal*, 30(4):298–307, 1987.