Bisimplicial separators

Martin Milanič[*] Irena Penev[†] Nevena Pivač[‡] Kristina Vušković[§]

March 18, 2022

Abstract

A *minimal separator* of a graph G is a set $S \subseteq V(G)$ such that there exist vertices $a, b \in V(G) \setminus S$ with the property that S separates a from b in G, but no proper subset of S does. For an integer $k \geq 0$, we say that a minimal separator is k-simplicial if it can be covered by k cliques and denote by $\mathcal{G}_k$ the class of all graphs in which each minimal separator is k-simplicial. We show that for each $k \geq 0$, the class $\mathcal{G}_k$ is closed under induced minors, and we use this to show that the Maximum Weight Stable Set problem can be solved in polynomial time for $\mathcal{G}_k$. We also give a complete list of minimal forbidden induced minors for $\mathcal{G}_2$. Next, we show that, for $k \geq 1$, every nonnull graph in $\mathcal{G}_k$ has a k-simplicial vertex, i.e., a vertex whose neighborhood is the union of k cliques; we deduce that the Maximum Weight Clique problem can be solved in polynomial time for graphs in $\mathcal{G}_2$. Further, we show that, for $k \geq 3$, it is NP-hard to recognize graphs in $\mathcal{G}_k$; the time complexity of recognizing graphs in $\mathcal{G}_2$ is unknown. We also show that the Maximum Clique problem is NP-hard for graphs in $\mathcal{G}_3$. Finally, we prove a decomposition theorem for diamond-free graphs in $\mathcal{G}_2$ (where the *diamond* is the graph obtained from K_4 by deleting one edge), and use this theorem to obtain polynomial-time algorithms for the Vertex Coloring and recognition problems for diamond-free graphs in $\mathcal{G}_2$, and improved running times for the Maximum Weight Clique and Maximum Weight Stable Set problems in this class of graphs.

Keywords: minimal separators, bisimplicial separators, induced minors, graph algorithms

1 Introduction

All graphs in this paper are finite, simple, and undirected. Our graphs may possibly be null. For a graph G and nonadjacent vertices $a, b \in V(G)$,

- an *(a, b)-separator* of G is a set $S \subseteq V(G) \setminus \{a, b\}$ such that a and b belong to distinct components of $G \setminus S$;
- a *minimal (a, b)-separator* of G is an (a, b)-separator S of G such that no proper subset of S is an (a, b)-separator of G.

[*]FAMNIT and IAM, University of Primorska, Koper, Slovenia. Email: martin.milanic@upr.si. Partially supported by the Slovenian Research Agency (I0-0035, research program P1-0285 and research projects J1-9110, N1-0102, N1-0160, J3-3001, J3-3002, and J3-3003).

[†]Computer Science Institute of Charles University (IÚUK), Prague, Czech Republic. Email: ipenev@iuuk.mff.cuni.cz.

[‡]FAMNIT and IAM, University of Primorska, Koper, Slovenia. Email: nevena.pivac@iam.upr.si. Partially supported by the Slovenian Research Agency (research program P1-0285, research projects N1-0102, J1-9110, and a Young Researchers Grant).

[§]School of Computing, University of Leeds, Leeds, UK. Email: k.vuskovic@leeds.ac.uk. Partially supported by EPSRC grant EP/V002813/1.

For a graph G, a set $S \subseteq V(G)$ is a *separator* (resp. *minimal separator*) of G if there exist distinct, nonadjacent vertices $a, b \in V(G) \setminus S$ such that S is an (a, b)-separator (resp. minimal (a, b)-separator) of G. Note that it is possible that S is a minimal separator of a graph G, even though some $S' \subsetneqq S$ is also a separator of G. Indeed, there may be a pair a, b of nonadjacent vertices such that S is a minimal (a, b)-separator of G, as well as some other pair a', b' of nonadjacent vertices such that some $S' \subsetneqq S$ is an (a', b')-separator of G.

A graph is *chordal* if it contains no induced cycles of length greater than three. Minimal separators have been studied since at least the 1960s, when chordal graphs were characterized as precisely those graphs in which all minimal separators are cliques [16]. Minimal separators were subsequently studied in [5] in the context of moplexes, have played an important role in sparse matrix computations via minimal triangulations (for a survey, see [19]), and have also had numerous algorithmic applications (see, e.g., [3, 7, 9, 24]). This paper is a contribution to the study of minimal separators.

For a class $\mathcal{C}$ of graphs, we denote by $\mathcal{G}_{\mathcal{C}}$ the class of all graphs G such that every minimal separator of G induces a graph from $\mathcal{C}$. Since complete graphs have no separators, we see for all classes $\mathcal{C}$, the class $\mathcal{G}_{\mathcal{C}}$ contains all complete graphs (including the null graph). For a nonnegative integer k, we denote by $\mathcal{G}_k$ the class of all graphs G that have the property that every minimal separator of G is the union of k (possibly empty) cliques. Obviously, $\mathcal{G}_0 \subseteq \mathcal{G}_1 \subseteq \mathcal{G}_2 \subseteq \dots$. Note that $\mathcal{G}_0$ is the class of all disjoint unions of complete graphs, and (by [16]) $\mathcal{G}_1$ is the class of all chordal graphs.

In this paper, we prove a number of results about classes of the form $\mathcal{G}_{\mathcal{C}}$, where $\mathcal{C}$ is a hereditary class. We place particular emphasis on the class $\mathcal{G}_2$. By the above, $\mathcal{G}_2$ contains all chordal graphs. Moreover, it is easy to see that all circular-arc graphs (that is, intersection graphs of arcs on a circle) belong to $\mathcal{G}_2$.

In Section 2, we prove some basic properties of the class $\mathcal{G}_{\mathcal{C}}$, when $\mathcal{C}$ satisfies various hypotheses.

In Section 3, we show that the classes $\mathcal{G}_k$ ($k \geq 0$) are closed under induced minors and describe the minimal forbidden induced minors in terms of minimally non-k-colorable graphs. In particular, this leads to a complete list of minimal forbidden induced minors for the class $\mathcal{G}_2$. Combining this with some results from the literature ([14, 18]), we show that for every integer $k \geq 0$, the Maximum Weight Stable Set Problem can be solved in polynomial time for graphs in $\mathcal{G}_k$, and we further show that all 1-perfectly-orientable graphs belong to $\mathcal{G}_2$.

In Section 4, we show that every nonnull graph in $\mathcal{G}_k$ ($k \geq 0$) has a k-simplicial vertex (i.e., a vertex whose neighborhood is the union of k cliques), and (using [25]) we deduce that the Maximum Weight Clique problem can be solved in polynomial time for graphs in $\mathcal{G}_2$.

In Section 5, we show that for each $k \geq 3$, it is NP-hard to recognize graphs in $\mathcal{G}_k$; the time complexity of recognizing graphs in $\mathcal{G}_2$ is unknown. We further show that the Maximum Clique problem is NP-hard for $\mathcal{G}_3$ (and consequently for $\mathcal{G}_k$ whenever $k \geq 3$). Note that, since Vertex Coloring is NP-hard for circular-arc graphs [17], which form a subclass of $\mathcal{G}_2$, the problem is also NP-hard for $\mathcal{G}_k$, whenever $k \geq 2$.

The *diamond* is the four-vertex graph obtained from the complete graph K_4 by deleting one edge. In Section 6, we prove a decomposition theorem for diamond-free graphs in $\mathcal{G}_2$, and use this theorem to obtain polynomial-time algorithms for the Vertex Coloring and recognition problems for diamond-free graphs in $\mathcal{G}_2$, and improved running times for the Maximum Weight Clique and Maximum Weight Stable Set problems in this class of graphs.

Table 1 summarizes our algorithmic and complexity results. Since $\mathcal{G}_0$ is the class of all disjoint unions of complete graphs, and $\mathcal{G}_1$ is the class of all chordal graphs, all problems from the table below can be solved in polynomial time for $\mathcal{G}_0$ and $\mathcal{G}_1$.

	diamond-free graphs in $\mathcal{G}_2$	$\mathcal{G}_2$	$\mathcal{G}_k$ $(k \geq 3)$
recognition	$\mathcal{O}(n(n+m))$	?	NP-hard
Maximum Weight Clique	$\mathcal{O}(n(n+m))$	$\mathcal{O}(n^4)$	NP-hard
Maximum Weight Stable Set	$\mathcal{O}(n^2(n+m))$	$\mathcal{O}(n^6)$	$\mathcal{O}(n^{2k+2})$
Vertex Coloring	$\mathcal{O}(n(n+m))$	NP-hard	NP-hard

Table 1: Summary of our algorithmic and complexity results. The number of vertices and edges of the input graph are denoted by n and m, respectively.

1.1 Terminology and notation

The vertex set and the edge set of a graph G are denoted by $V(G)$ and $E(G)$, respectively. The complement of G is denoted by $\overline{G}$.

A *clique* in a graph G is a (possibly empty) set of pairwise adajcent vertices, and a *stable set* in G is a (possibly empty) set of pairwise nonadjacent vertices. Given a graph G and a vertex $x \in V(G)$, we denote by $N_G(x)$ the set of all neighbors of x in G, and we set $N_G[x] := \{x\} \cup N_G(x)$. For a nonnegative integer k, a vertex $x \in V(G)$ is *k-simplicial* in G if $N_G(x)$ is the union of k (possibly empty) cliques. A 1-simplicial vertex is also called *simplicial*, and a 2-simplicial vertex is also called *bisimplicial*. Analogously, for a graph G and a set $X \subseteq V(G)$, we say that X is *k-simplicial* in G if it is the union of k (possibly empty) cliques and *bisimplicial* if it is 2-simplicial. In particular, a graph G belongs to the class $\mathcal{G}_k$ if and only if every minimal separator of G is k-simplicial, and to $\mathcal{G}_2$ if and only if every minimal separator of G is bisimplicial.

For a graph G and a set $X \subseteq V(G)$, $G[X]$ is the subgraph of G induced by X; if $X = \{x_1, \ldots, x_t\}$, we sometimes write $G[x_1, \ldots, x_t]$ instead of $G[X]$. Furthermore, $G \setminus X$ is the graph obtained from G by deleting all vertices in X, i.e., $G \setminus X := G[V(G) \setminus X]$; if $X = \{x\}$, we sometimes write $G \setminus x$ instead of $G \setminus X$. A *path* in G is a nonempty sequence $p_1, \ldots, p_k$ of pairwise distinct vertices of G such that p_i and p_{i+1} are adjacent for all $i \in \{1, \ldots, k-1\}$.

For a nonnegative integer n, we denote by K_n, P_n, and C_n, respectively, the complete graph, path graph, and cycle graph on n vertices. For two nonnegative integers p and q, we denote by $K_{p,q}$ the complete bipartite graph with parts of size p and q. A class $\mathcal{C}$ of graphs is *hereditary* if for all $G \in \mathcal{C}$, all (isomorphic copies of) induced subgraphs of G belong to $\mathcal{C}$.

Given a graph G and an edge $e = xy \in E(G)$, *subdividing* the edge e means replacing e with a path of length two, that is, removing the edge e and adding a new vertex z adjacent to precisely x and y. A *subdivision* of a graph G is any graph obtained from G by a (possibly null) sequence of edge subdivisions. Given a graph G and an edge $e \in E(G)$, *contracting* the edge $e = xy$ means replacing the vertices x and y in G with a new vertex v^{xy} adjacent to every vertex that is adjacent in G to x or y. We denote by G/e the graph obtained from G by contracting e. We say that a graph H is an *induced minor* of a graph G if H can be obtained from G via a (possibly null) sequence of vertex deletions and edge contractions. Given graphs H and G, we say that G is *H-induced-minor-free* if H is not an induced minor of G. For a family $\mathcal{H}$ of graphs, we say that a graph G is *$\mathcal{H}$-induced-minor-free* if G is H-induced-minor-free for all graphs $H \in \mathcal{H}$.

Given a graph G, a vertex $x \in V(G)$, and a set $Y \subseteq V(G) \setminus \{x\}$, we say that x is *complete* (resp. *anticomplete*) to Y in G if x is adjacent (resp. nonadjacent) to all vertices in Y. Given disjoint sets $A, B \subseteq V(G)$, we say that A is *complete* (resp. *anticomplete*) to B in G if every vertex in A is complete (resp. anticomplete) to B in G.

A *cutset* of a graph G is a (possibly empty) set $C \subseteq V(G)$ such that $G \setminus C$ is disconnected.

A *minimal cutset* of a graph G is a cutset C of G such that no proper subset of C is a cutset of G. A *clique cutset* is a cutset that is a clique. Note that $\emptyset$ is a clique cutset of any disconnected graph.

A *cut partition* of a graph G is a partition (A, B, C) of $V(G)$ such that A and B are nonempty and anticomplete to each other (the set C may possibly be empty). Note that if (A, B, C) is a cut partition of G, then C is a cutset of G; conversely, every cutset of G gives rise to at least one cut partition of G.

If H_1 and H_2 are graphs on disjoint vertex sets, the *disjoint union* of H_1 and H_2, denoted by $H_1 \cup H_2$, is the graph with vertex set $V(H_1) \cup V(H_2)$ and edge set $E(H_1) \cup E(H_2)$.

Given a graph $G = (V, E)$ and a vertex weight function $w : V \to \mathbb{Q}_+$, the Maximum Weight Clique problem is the problem of computing a clique C in G with maximum total weight, where the weight of a set $X \subseteq V$ is defined as $\sum_{x \in X} w(x)$. Similarly, the Maximum Weight Stable Set problem is the problem of computing a stable set S in G with maximum total weight. In the case of weight functions constantly equal to 1, we obtain the Maximum Clique and Maximum Stable Set problems, respectively. A graph $G = (V, E)$ is *k-colorable* if V is the union of k stable sets in G. The Vertex Coloring is the following problem: given a graph G, compute its *chromatic number* $\chi(G)$, that is, the smallest integer k such that G is k-colorable. For a positive integer k, the k-Coloring is the following decision problem: given a graph G, determine whether G is k-colorable.

2 Basic properties

Recall that for any graph class $\mathcal{C}$, we denote by $\mathcal{G}_{\mathcal{C}}$ the class of all graphs G such that every minimal separator of G induces a graph from $\mathcal{C}$. The following proposition provides some equivalent characterizations of the class $\mathcal{G}_{\mathcal{C}}$ for the case when $\mathcal{C}$ is hereditary.

Proposition 2.1. *Let $\mathcal{C}$ be a hereditary class of graphs, and let G be a graph. Then the following are equivalent:*

(a) $G \in \mathcal{G}_{\mathcal{C}}$;

(b) for all induced subgraphs H of G, every minimal separator S of H satisfies $H[S] \in \mathcal{C}$;

(c) for all induced subgraphs H of G, every minimal cutset C of H satisfies $H[C] \in \mathcal{C}$.

Proof. We prove the result by showing that (a) implies (b), that (b) implies (c), and that (c) implies (a).

First, we assume that (a) holds, and we prove (b). Let H be an induced subgraph of G, and suppose that $a, b \in V(H)$ are distinct, nonadjacent vertices, and that S is a minimal (a, b)-separator of H. Then $S \cup (V(G) \setminus V(H))$ is an (a, b)-separator of G. Let $S^* \subseteq S \cup (V(G) \setminus V(H))$ be a minimal (a, b)-separator of G; by (a), $G[S^*] \in \mathcal{C}$. Since S^* is an (a, b)-separator of G, we have that $S^* \cap V(H)$ is an (a, b)-separator of H. Moreover, $S^* \cap V(H) \subseteq S$, and so the minimality of S guarantees that $S^* \cap V(H) = S$; consequently, $S \subseteq S^*$. Since $G[S^*] \in \mathcal{C}$, and since $\mathcal{C}$ is hereditary, it follows that $G[S] \in \mathcal{C}$. Clearly, $G[S] = H[S]$, and so $H[S] \in \mathcal{C}$. Thus, (b) holds.

Next, we assume that (b) holds, and we prove (c). Let H be an induced subgraph of G, and suppose that C is a minimal cutset of H. Let A and B be the vertex sets of two distinct components of $H \setminus C$. The minimality of C guarantees that every vertex in C has a neighbor both in A and in B, and this, in turn, guarantees that for all $a \in A$ and $b \in B$, C is a minimal (a, b)-separator of H. But now (b) implies that $H[C] \in \mathcal{C}$. Thus, (c) holds.

Finally, we assume that (c) holds, and we prove (a). Suppose that $a, b \in V(G)$ are distinct, nonadjacent vertices, and that S is a minimal (a, b)-separator of G. Let A (resp. B) be the vertex set of the component of $G \setminus S$ that contains a (resp. b). Clearly, A and B are disjoint and anticomplete to each other. Furthermore, the minimality of S implies that every vertex in S has a neighbor both in A and in B. Set $H := G[A \cup B \cup S]$. Then (A, B, S) is a cut partition of H; furthermore, since $H[A]$ and $H[B]$ are connected, and every vertex of S has a neighbor both in A and in B, we see that S is a minimal cutset of H. Now (c) guarantees that $H[S] \in \mathcal{C}$; since $H[S] = G[S]$, it follows that $G[S] \in \mathcal{C}$. Thus, (a) holds. □

Corollary 2.2. *Let $\mathcal{C}$ be a hereditary class of graphs. Then $\mathcal{G}_{\mathcal{C}}$ is hereditary.*

Proof. This readily follows from Proposition 2.1, and more precisely, from the equivalence of (a) and (b) from Proposition 2.1. □

Theorem 2.3. *Let $\mathcal{C}$ be a hereditary class of graphs that contains all complete graphs.*[1] *Let G be a graph that admits a clique cutset C and let (A, B, C) be an associated cut partition of G. Assume that $G_A := G[A \cup C]$ and $G_B := G[B \cup C]$ both belong to $\mathcal{G}_{\mathcal{C}}$. Then $G \in \mathcal{G}_{\mathcal{C}}$.*

Proof. Fix a pair of distinct, nonadjacent vertices $x, y \in V(G)$, and let S be a minimal (x, y)-separator of G. We must show that $G[S] \in \mathcal{C}$. Since C is a clique, it contains at most one of x, y; by symmetry, we may therefore assume that $x \in A$. We now consider two cases: when $y \in A \cup C$, and when $y \in B$.

Case 1: $y \in A \cup C$. Then $S \cap (A \cup C)$ is an (x, y)-separator of G_A; let $S' \subseteq S \cap (A \cup C)$ be a minimal (x, y)-separator of G_A. Since $G_A \in \mathcal{G}_{\mathcal{C}}$, we see that $G_A[S'] = G[S']$ belongs to $\mathcal{C}$. If $S' = S$, then we are done. So, assume that $S' \subsetneqq S$. The minimality of S then implies that there is a path $p_1, \dots, p_s$ in $G \setminus S'$, with $p_1 = x$ and $p_s = y$. Since S' is an (x, y)-separator of G_A, we see that at least one vertex of the path $p_1, \dots, p_s$ belongs to B. Let i be the smallest index in $\{1, \dots, s\}$ such that $p_i \in B$, and let j be the largest index in $\{1, \dots, s\}$ such that $p_j \in B$. Since $p_1, p_s \in A \cup C$, we have that $2 \leq i \leq j \leq s - 1$. Moreover, since A is anticomplete to B, we have that $p_{i-1}, p_{j+1} \in C$; since C is a clique, we see that p_{i-1}, p_{j+1} are adjacent. But now $p_1, \dots, p_{i-1}, p_{j+1}, \dots, p_s$ is a path between x and y in $G_A \setminus S'$, contrary to the fact that S' is an (x, y)-separator of G_A.

Case 2: $y \in B$. Note that, in this case, C is an (x, y)-separator of G.

Claim. At least one of $S \cap A$ and $S \cap B$ is empty.

Proof of the Claim. Suppose otherwise, i.e., that both $S \cap A$ and $S \cap B$ are nonempty. Set $S_A := S \setminus B$ and $S_B := S \setminus A$. By the minimality of S, there is a path $p_1, \dots, p_s$, with $p_1 = x$ and $p_s = y$, in $G \setminus S_A$, and there is a path $q_1, \dots, q_t$, with $q_1 = x$ and $q_t = y$, in $G \setminus S_B$. Since $p_1 \in A$ and $p_s \in B$, and since A is anticomplete to B, some internal vertex of the path $p_1, \dots, p_s$ belongs to C; let i be the smallest index in $\{2, \dots, s-1\}$ such that $p_i \in C$ (then $p_1, \dots, p_{i-1} \in A$). Similarly, at least one internal vertex of the path $q_1, \dots, q_t$ belongs to C; let j be the largest index in $\{2, \dots, t-1\}$ such that $q_j \in C$ (then $q_{j+1}, \dots, q_t \in B$). Since C is a clique, we see that p_i and q_j are either equal or adjacent. In the former case, $p_1, \dots, p_i, q_{j+1}, \dots, q_t$ is a path between x and y in $G \setminus S$; and in the latter case, $p_1, \dots, p_i, q_j, \dots, q_t$ is a path between x and y in $G \setminus S$. But neither outcome is possible, since S is an (x, y)-separator of G. This proves the Claim. ♦

By the Claim, and by symmetry, we may assume that $S \cap B = \emptyset$, i.e., $S \subseteq A \cup C$. Let Y be the vertex set of the component of $G[B]$ that contains y.

Suppose first that $C \setminus S$ is anticomplete to Y. Then $C \cap S$ is an (x, y)-separator of G, and so the minimality of S guarantees that $S \subseteq C$. Thus, S is a clique; since $\mathcal{C}$ contains all complete graphs, it follows that $G[S] \in \mathcal{C}$, and we are done.

[1] However, not all graphs in $\mathcal{C}$ need be complete.

From now on, we assume that $C \setminus S$ is not anticomplete to Y. Fix a vertex $c \in C \setminus S$ that has a neighbor in Y. Since $y \in Y$, and $G[Y]$ is connected, the graph G contains a path $q_1, \dots, q_t$, with $q_1 = c$, $q_t = y$, and $q_2, \dots, q_t \in Y$ (so, $q_2, \dots, q_t \in B$). Now, suppose that there is a path $p_1, \dots, p_s$ in $G_A \setminus S$, with $p_1 = x$ and $p_s = c$. Then $p_1, \dots, p_s, q_2, \dots, q_t$ is a path in $G \setminus S$ between x and y, contrary to the fact that S is an (x, y)-separator of G. So, S is an (x, c)-separator of G_A. Let $S' \subseteq S$ be a minimal (x, c)-separator of G_A; since $G_A \in \mathcal{G}_{\mathcal{C}}$, we see that $G_A[S'] = G[S']$ belongs to $\mathcal{C}$. If S' is an (x, y)-separator of G, then the minimality of S guarantees that $S = S'$, and we are done. So, assume that S' is not an (x, y)-separator of G. Then there is a path $r_1, \dots, r_k$ in $G \setminus S'$, with $r_1 = x$ and $r_k = y$. Since $x \in A$, $y \in B$, and A is anticomplete to B, we see that some internal vertex of $r_1, \dots, r_k$ belongs to C; let i be the smallest index in $\{2, \dots, k-1\}$ such that $r_i \in C$. Since $r_i, c \in C$, and C is a clique, we see that r_i and c are either equal or adjacent. In the former case, $r_1, \dots, r_i$ is a path from x to c in $G_A \setminus S'$, and in the latter case, $r_1, \dots, r_i, c$ is a path from x to c in $G_A \setminus S'$. But neither outcome is possible, since S' is an (x, c)-separator of G_A. □

Corollary 2.4. *Let k be a positive integer, let G be a graph that admits a clique cutset C, and let (A, B, C) be an associated cut partition of G. Assume that $G[A \cup C]$ and $G[B \cup C]$ both belong to $\mathcal{G}_k$. Then $G \in \mathcal{G}_k$.*

Proof. This follows immediately from Theorem 2.3. □

Note that Corollary 2.4 fails for $\mathcal{G}_0$: the two-edge path P_3 is an obvious counterexample.

3 A forbidden induced minor characterization of $\mathcal{G}_2$

In this section we show that the classes $\mathcal{G}_k$, $k \geq 0$, are closed under induced minors, describe the minimal forbidden induced minors, and examine some algorithmic consequences of these results.

The fact that the classes $\mathcal{G}_k$ are closed under induced minors is a consequence of the following more general result.

Theorem 3.1. *Let $\mathcal{C}$ be a hereditary class of graphs, closed under edge addition. Then the class $\mathcal{G}_{\mathcal{C}}$ is closed under induced minors.*

Proof. By Corollary 2.2, $\mathcal{G}_{\mathcal{C}}$ is hereditary. So, it suffices to show that $\mathcal{G}_{\mathcal{C}}$ is closed under edge contractions. Fix $G \in \mathcal{G}_{\mathcal{C}}$, let xy be an edge of G, and set $G' := G/xy$; the vertex of G' to which the edge xy is contracted will be denoted by v^{xy}. We must show that $G' \in \mathcal{G}_{\mathcal{C}}$, i.e., that for any minimal separator S of G', we have that $G'[S] \in \mathcal{C}$.

We first deal with minimal separators of G' that contain v^{xy}. So, suppose that $S \subseteq V(G')$ is a minimal separator of G' such that $v^{xy} \in S$; we must show that $G'[S] \in \mathcal{C}$. Fix distinct $a, b \in V(G') \setminus S$ such that S is a minimal (a, b)-separator of G'. Then $S^* := (S \setminus \{v^{x,y}\}) \cup \{x, y\}$ is an (a, b)-separator of G. Let $S' \subseteq S^*$ be a minimal (a, b)-separator of G; since $G \in \mathcal{G}_{\mathcal{C}}$, we have that $G[S'] \in \mathcal{C}$. If $x, y \notin S'$, then $S' \subsetneq S$ is an (a, b)-separator of G', contrary to the minimality of S. So, S' contains at least one of x, y. Then $(S' \setminus \{x, y\}) \cup \{v^{xy}\}$ is an (a, b)-separator of G'; since $(S' \setminus \{x, y\}) \cup \{v^{xy}\} \subseteq S$, the minimality of S implies that $S = (S' \setminus \{x, y\}) \cup \{v^{xy}\}$. As we show next, the graph $G'[S]$ can be obtained from an induced subgraph of $G[S']$ by possibly adding some edges. By symmetry, we may assume without loss of generality that $x \in S'$. Since $S \setminus \{v^{xy}\} = S' \setminus \{x, y\}$, the graph $G'[S]$ is isomorphic to the graph obtained from the subgraph of $G[S']$ induced by $S' \setminus \{y\}$ by adding to it the edges from x to all vertices in $S' \setminus \{x, y\}$ that are adjacent in G to y but not to x. Since $G[S'] \in \mathcal{C}$, and $\mathcal{C}$ is hereditary and closed under edge addition, we deduce that $G'[S] \in \mathcal{C}$, and we are done.

It remains to consider minimal separators of G' that do not contain v^{xy}. Here, we first observe that for any pair of nonadjacent vertices a, b of G', and any set $S \subseteq V(G') \setminus \{a, b, v^{xy}\} = V(G) \setminus \{a, b, x, y\}$, both the following hold:

(1) if $v^{xy} \notin \{a, b\}$, then S is an (a, b)-separator of G if and only if S is an (a, b)-separator of G';

(2) if $v^{xy} = a$, then S is an (x, b)-separator of G if and only if S is an (a, b)-separator of G'.

Clearly, (1) and (2) imply that any set $S \subseteq V(G') \setminus \{v^{xy}\} = V(G) \setminus \{x, y\}$ is a minimal separator of G' if and only if it is a minimal separator of G. But for any $S \subseteq V(G') \setminus \{v^{xy}\} = V(G) \setminus \{x, y\}$, we have that $G'[S] = G[S]$, and moreover, if S is a minimal separator of G, then $G[S] \in \mathcal{C}$. This shows that if a set $S \subseteq V(G') \setminus \{v^{xy}\}$ is a minimal separator of G', then $G'[S] \in \mathcal{C}$. This completes the proof. □

Corollary 3.2. *For all integers $k \geq 0$, the class $\mathcal{G}_k$ is closed under induced minors.*

Proof. For $k \geq 1$, this follows immediately from Theorem 3.1. For $k = 0$, we observe that $\mathcal{G}_0$ is the class of all graphs in which all minimal separators are empty, i.e., $\mathcal{G}_0$ is precisely the class of all disjoint unions of (arbitrarily many) complete graphs. But now it is obvious that $\mathcal{G}_0$ is closed under induced minors. □

The reader may wonder whether, in Theorem 3.1, it is necessary to assume that $\mathcal{C}$ is closed under edge addition. Here, we note that if $\mathcal{C}$ is the class of all edgeless graphs, then $K_{2,3} \in \mathcal{G}_{\mathcal{C}}$, but contracting one edge of $K_{2,3}$ produces the diamond, which does not belong to $\mathcal{G}_{\mathcal{C}}$. Thus, the assumption about edge additions cannot be simply removed from Theorem 3.1, although it is possible that some other (weaker) assumption would suffice instead.

A graph H is a *forbidden induced minor* for a class $\mathcal{G}$ if every graph in $\mathcal{G}$ is H-induced-minor-free. A forbidden induced minor H for a class $\mathcal{G}$ is *minimal* if every proper induced minor of H is an induced minor of some graph in $\mathcal{G}$. Clearly, a class $\mathcal{G}$ is closed under induced minors if and only if it is precisely the class of all $\mathcal{M}$-induced-minor-free graphs, where $\mathcal{M}$ is the collection of all minimal forbidden induced minors for $\mathcal{G}$. For each integer $k \geq 0$, we denote by $\mathcal{M}_k$ the class of all minimal forbidden induced minors for $\mathcal{G}_k$; by Corollary 3.2, the class $\mathcal{G}_k$ is closed under induced minors, and so $\mathcal{G}_k$ is precisely the class of all $\mathcal{M}_k$-induced-minor-free graphs.

The class $\mathcal{G}_0$ is precisely the class of all disjoint unions of (arbitrarily many) complete graphs; consequently, $\mathcal{M}_0 = \{P_3\}$. Graphs in $\mathcal{G}_1$ are precisely the graphs in which all minimal separators are cliques. Thus (see [16]), $\mathcal{G}_1$ is precisely the class of *chordal* graphs, i.e., graphs that contain no induced cycles of length greater than three, or equivalently, graphs that are C_4-induced-minor-free. So, $\mathcal{M}_1 = \{C_4\}$. Note that the graphs P_3 and C_4 are isomorphic, respectively, to the graphs $\overline{K_2 \cup K_1}$ and $\overline{K_2 \cup K_2}$. In the remainder of this section, we prove that $\mathcal{M}_2 = \{\overline{K_2 \cup C_{2k+1}} \mid k \in \mathbb{N}\}$ (see Fig. 1), as a consequence of a more general description of the set of forbidden induced minors for the class $\mathcal{G}_k$. For an integer $k \geq 0$, let $\mathcal{F}_k$ denote the class of graphs G such that G is not k-colorable but every proper induced subgraph of G is k-colorable. For instance, $\mathcal{F}_0 = \{K_1\}$, $\mathcal{F}_1 = \{K_2\}$, and $\mathcal{F}_2 = \{C_{2k+1} \mid k \in \mathbb{N}\}$.

Theorem 3.3. *For every integer $k \geq 0$, we have $\mathcal{M}_k = \{\overline{K_2 \cup F} \mid F \in \mathcal{F}_k\}$.*

Proof. To simplify notation, we set $\mathcal{M} := \{\overline{K_2 \cup F} \mid F \in \mathcal{F}_k\}$. We must show that $\mathcal{M}_k = \mathcal{M}$. First we show that graphs in $\mathcal{M}$ are pairwise incomparable with respect to the induced minor relation. Consider two distinct graphs $H, H' \in \mathcal{M}$ and suppose for a contradiction that H is an induced minor of H'. This means that the graph $\overline{H}$ can be obtained from an induced subgraph of the graph $\overline{H'}$ by a sequence of *cocontracting non-edges*, where the operation of cocontracting

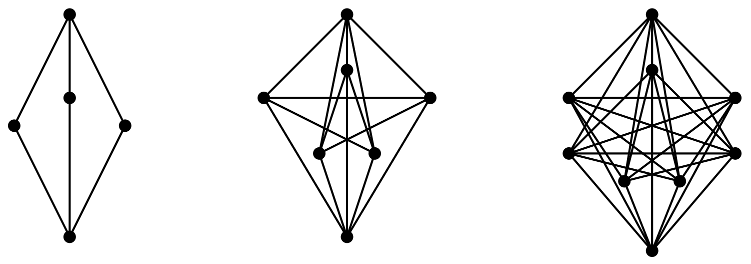

Figure 1: Some small graphs in $\mathcal{M}_2$.

a non-edge xy in a graph G means replacing the pair of non-adjacent vertices x and y with a new vertex v^{xy} adjacent precisely to the common neighbors of x and y in G. (Equivalently, we take the complement of the graph obtained from the graph $\overline{G}$ by contracting the edge xy.) By definition of $\mathcal{M}_k$, there exists graphs $F, F' \in \mathcal{F}_k$ such that $H = \overline{K_2 \cup F}$ and $H' = \overline{K_2 \cup F'}$. Then $\overline{H} = K_2 \cup F$ and $\overline{H'} = K_2 \cup F'$. Since $K_2 \cup F$ can be obtained from an induced subgraph of the graph $K_2 \cup F'$ by a sequence of cocontracting non-edges, it follows in particular that the graph $K_2 \cup F$ is a subgraph of $K_2 \cup F'$. Since F and F' are distinct graphs from $\mathcal{F}_k$, none of them is a subgraph of the other. On the other hand, since F is a connected subgraph of $K_2 \cup F'$ but not of F', we infer that F is a subgraph of K_2 and consequently that the connected graph F' is isomorphic to K_1. It follows that $K_2 \cup F'$ is isomorphic to $K_2 \cup K_1$. However, since the graph F has at least one vertex, the fact that the graph $K_2 \cup F$ is a proper subgraph of $K_2 \cup F' \cong K_2 \cup K_1$ leads to a contradiction. This shows that the graphs in $\mathcal{M}$ are pairwise incomparable with respect to the induced minor relation.

It remains to show that $\mathcal{G}_k$ is precisely the class of all $\mathcal{M}$-induced-minor-free graphs.

Let us first show that all graphs in $\mathcal{G}_k$ are $\mathcal{M}$-induced-minor-free. Since $\mathcal{G}_k$ is closed under induced minors (by Corollary 3.2), it suffices to show that no graph in $\mathcal{M}$ belongs to $\mathcal{G}_k$. Fix $M \in \mathcal{M}$. Then M contains two nonadjacent vertices, a and b, such that $\{a, b\}$ is complete to $V(M) \setminus \{a, b\}$ in M, and $M \setminus \{a, b\}$ is the complement of a graph $F \in \mathcal{F}_k$. But now $V(M) \setminus \{a, b\}$ is a minimal (a, b)-separator of M, and, since F is not k-colorable, $V(M) \setminus \{a, b\}$ is not a union of k cliques in M. So, $M \notin \mathcal{G}_k$.

For the other direction, suppose that G is $\mathcal{M}$-induced-minor-free. We need to show that $G \in \mathcal{G}_k$. Suppose otherwise; then there exist two nonadjacent vertices a and b in G such that some minimal (a, b)-separator S of G is not a union of k cliques. Then $\overline{G}[S]$ is not k-colorable, and so there exists a subset $X \subseteq S$ such that $\overline{G}[X] \in \mathcal{F}_k$. Let C_a (resp. C_b) denote the component of $G \setminus S$ containing a (resp. b). Let $G' := G[V(C_a) \cup V(C_b) \cup X]$, and let G'' be the graph obtained from G' by contracting all the edges in C_a and in C_b. The components C_a and C_b get contracted into two nonadjacent vertices, say a^* and b^*, respectively. Moreover, since S is a minimal (a, b)-separator of G, every vertex in $X \subseteq S$ has a neighbor (in G) in both C_a and in C_b. Therefore, a^* and b^* are both complete to X in G'', and we conclude that $G'' \in \mathcal{M}$. But by construction, G'' is an induced minor of G, contrary to the assumption that G is $\mathcal{M}$-induced-minor-free. □

For the case $k = 2$, Theorem 3.3 implies the following explicit characterization of the set of forbidden induced minors for the class $\mathcal{G}_2$.

Corollary 3.4. *$\mathcal{M}_2 = \{\overline{K_2 \cup C_{2k+1}} \mid k \in \mathbb{N}\}$.*

A graph is 1-*perfectly-orientable* if it admits an orientation in which the out-neighborhood of each vertex is a clique of the underlying graph. It was shown in [18] that all 1-perfectly-orientable graphs are $\{\overline{K_2 \cup C_{2k+1}} \mid k \in \mathbb{N}\}$-induced-minor-free. Thus, Corollary 3.4 implies that 1-perfectly-orientable graphs form a subclass of $\mathcal{G}_2$.

3.1 Algorithmic considerations

It was shown by Dallard et al. in [14] that, for each positive integer k, the Maximum Weight Stable Set problem can be solved in $\mathcal{O}(n^{2k})$ time for n-vertex $K_{2,k}$-induced-minor-free graphs. To connect this result with the classes $\mathcal{G}_k$, note that for each integer $k \geq 0$, the complete graph with $k+1$ vertices belongs to $\mathcal{F}_k$; hence, since the graph $\overline{K_2 \cup K_{k+1}}$ is isomorphic to the complete bipartite graph $K_{2,k+1}$, Theorem 3.3 implies that every graph in $\mathcal{G}_k$ is $K_{2,k+1}$-induced-minor-free. Therefore, the result by Dallard et al. implies the following.

Corollary 3.5. *For each integer $k \geq 0$, the* Maximum Weight Stable Set *problem can be solved in $\mathcal{O}(n^{2k+2})$ time for n-vertex graphs in $\mathcal{G}_k$.*

Similar results hold for a number of other related problems, including the Maximum Induced Matching problem, the Dissociation Set problem, etc. We refer to [14] for details.

Furthermore, Dallard et al. showed in [15] that in any class of $K_{2,k}$-induced-minor-free graphs, the treewidth of the graphs in the class is bounded from above by some polynomial function of the clique number (see also [14]). Combining this with a result of Chaplick et al. [11, Theorem 12], it follows that for any two positive integers k and ℓ, the ℓ-Coloring problem is solvable in time $\mathcal{O}(n)$ in the class of n-vertex $K_{2,k}$-induced-minor-free graphs, and thus in the class $\mathcal{G}_k$ as well. (The $\mathcal{O}$-notation hides a constant depending on k and ℓ.) The same result holds in fact for the more general List ℓ-Coloring problem, in which every vertex is equipped with a list of available colors from the set $\{1, \ldots, \ell\}$.

4 Vertex neighborhoods

Recall that a vertex v in a graph G is *k-simplicial* it its neighborhood is a union of k cliques. In this section we show that for every $k \geq 0$, every nonnull graph in $\mathcal{G}_k$ has a k-simplicial vertex and examine some algorithmic consequences of this result for the case $k = 2$.

For a family $\mathcal{F}$ of graphs, an ordering $v_1, \ldots, v_n$ of the vertices of a graph G is an *$\mathcal{F}$-elimination ordering* if for every index $i \in \{1, \ldots, n\}$, the graph $G[N_{G[v_1,\ldots,v_i]}(v_i)]$ is $\mathcal{F}$-free. Note that a graph G admits an $\mathcal{F}$-elimination ordering if and only if every nonnull induced subgraph of G contains a vertex whose neighborhood induces an $\mathcal{F}$-free subgraph in G.

LexBFS is a linear time algorithm of Rose, Tarjan, and Lueker [23] whose input is any nonnull graph G together with a vertex $s \in V(G)$, and whose output is an ordering of the vertices of G starting at s. An ordering of the vertices of a graph G is a *LexBFS ordering* if there exists a vertex s of G such that the ordering can be produced by LexBFS when the input is G, s.

In certain cases, $\mathcal{F}$-elimination orderings can be found using LexBFS. This relies on the concept of locally $\mathcal{F}$-decomposable graphs and graph classes, introduced by Aboulker et al. in [1]. Let $\mathcal{F}$ be a family of graphs. A graph G is *locally $\mathcal{F}$-decomposable* if for every vertex v of G, every $F \in \mathcal{F}$ contained in $G[N_G(v)]$ and every component C of $G \setminus N_G[v]$, there exists $y \in V(F)$ such that y has a nonneighbor in F and has no neighbors in C. A class of graphs $\mathcal{G}$ is *locally $\mathcal{F}$-decomposable* if every graph $G \in \mathcal{G}$ is a locally $\mathcal{F}$-decomposable graph.

Theorem 4.1 (Aboulker et al. [1])**.** *If $\mathcal{F}$ is a family of noncomplete graphs and G is locally $\mathcal{F}$-decomposable, then every LexBFS ordering of G is an $\mathcal{F}$-elimination ordering.*

For a hereditary graph class $\mathcal{C}$, we denote by $\mathcal{F}_{\mathcal{C}}$ the class of all graphs G such that G does not belong to $\mathcal{C}$, but all proper induced subgraphs of G do belong to $\mathcal{C}$. A *universal vertex* of a graph G is a vertex u such that $N_G[u] = V(G)$. *Adding a universal vertex* to a graph G means adding a new vertex v and making it adjacent to all vertices of G; note that v is a universal vertex in the resulting graph.

Theorem 4.2. *Let $\mathcal{C}$ be a hereditary class of graphs that is closed under the addition of universal vertices. Then every nonnull graph $G \in \mathcal{G}_{\mathcal{C}}$ contains a vertex v such that $G[N_G(v)] \in \mathcal{C}$.*

Proof. First, we note that no graph in $\mathcal{F}_{\mathcal{C}}$ contains a universal vertex. Indeed, if some $F \in \mathcal{F}_{\mathcal{C}}$ contained a universal vertex u, then the definition of $\mathcal{F}_{\mathcal{C}}$ would imply that $F \setminus u$ belongs to $\mathcal{C}$, and since $\mathcal{C}$ is closed under the addition of universal vertices, it would follow that $F \in \mathcal{C}$, a contradiction.

Now, fix a nonnull graph $G \in \mathcal{G}_{\mathcal{C}}$ and a vertex $x \in V(G)$. Suppose that F is an induced subgraph of $G[N_G(x)]$ such that $F \in \mathcal{F}_{\mathcal{C}}$. By the above, F does not contain a universal vertex, and consequently, every vertex of F has a nonneighbor in F. Let C be a component of $G \setminus N_G[x]$; we must show that some vertex in F is anticomplete to $V(C)$. Suppose otherwise, that is, suppose that every vertex in $V(F)$ has a neighbor in $V(C)$. Let $z \in V(C)$. Clearly, $N_G(x)$ is an (x, z)-separator of G, and moreover, any minimal (x, z)-separator of G included in $N_G(x)$ includes $V(F)$; since $G \in \mathcal{G}_{\mathcal{C}}$ and $\mathcal{C}$ is hereditary, it follows that $F \in \mathcal{C}$, contrary to the fact that $F \in \mathcal{F}_{\mathcal{C}}$. Thus, some vertex in $V(F)$ is indeed anticomplete to $V(C)$. It follows that G is locally $\mathcal{F}_{\mathcal{C}}$-decomposable, and so by Theorem 4.1, G contains a vertex whose neighborhood induces an $\mathcal{F}_{\mathcal{C}}$-free subgraph. But clearly, every $\mathcal{F}_{\mathcal{C}}$-free graph belongs to $\mathcal{C}$. This completes the proof. □

The reader may wonder whether, in Theorem 4.2, it might be possible to eliminate the hypothesis that $\mathcal{C}$ is closed under the addition of universal vertices. This would in fact not be possible (at least not without adding some other, perhaps weaker, hypothesis). To see this, fix any positive integer ℓ, and any hereditary class $\mathcal{C}$ that does not contain K_ℓ. The class $\mathcal{G}_{\mathcal{C}}$ contains all complete graphs, and in particular, $K_{\ell+1} \in \mathcal{G}_{\mathcal{C}}$. However, the neighborhood of any vertex of $K_{\ell+1}$ induces a K_ℓ, and $K_\ell \notin \mathcal{C}$.

Corollary 4.3. *Let k be a positive integer. Then every nonnull graph in $\mathcal{G}_k$ contains a k-simplicial vertex.*

Proof. To derive this result from Theorem 4.2, it suffices to show that for for all $k \geq 1$, the class $\mathcal{G}_k$ is closed under the addition of universal vertices. Let G be an arbitrary graph in $\mathcal{G}_k$ and let G' be the graph obtained from G by adding to it a universal vertex v. Fix two distinct, nonadjacent vertices $x, y \in V(G')$ and let S be a minimal (x, y)-separator of G'. Then $x, y \in V(G)$ and consequently, $v \in S$ and $S \setminus \{v\}$ is a minimal separator of G. Since G belongs to $\mathcal{G}_k$, it contains k cliques $C_1, \ldots, C_k$ with union $S \setminus \{v\}$. But then $C_1, \ldots, C_k \cup \{v\}$ are k cliques in G' covering the minimal separator S. It follows that $G' \in \mathcal{G}_k$ and hence the class $\mathcal{G}_k$ is closed under the addition of universal vertices. □

Remark 4.4. Corollary 4.3 can also be obtained from the fact that every nonnull graph G contains a *moplex* [5], that is, a clique C such that every two vertices in C have the same closed neighborhood and the neighborhood of C is either empty or a minimal separator of G. (In fact, the last vertex visited by any execution of LexBFS on G necessarily belongs to a moplex.) Given a graph $G \in \mathcal{G}_k$ and a vertex v that belongs to a moplex C of G, there exist k cliques $C_1, \ldots, C_k$ covering the neighborhood of C. But then $C_1, \ldots, C_k \cup (C \setminus \{v\})$ are k cliques covering $N_G(v)$, showing that v is a k-simplicial vertex.

4.1 Algorithmic implications

A *bisimplicial elimination ordering* of a graph G is an ordering $v_1, \dots, v_n$ of the vertices of G such that for all $i \in \{1, \dots, n\}$, v_i is bisimplicial in the graph $G[v_1, \dots, v_i]$. The Maximum Weight Clique problem can be solved in time $\mathcal{O}(n^4)$ for n-vertex graphs that admit a bisimplicial elimination ordering, see [25]. The algorithm iteratively removes bisimplicial vertices and reduces the problem to solving n instances of the Maximum Weight Stable Set problem in bipartite graphs. The polynomial running time of the algorithm given in [25] was based on polynomial-time solvability of the Maximum Weight Stable Set problem in the class of perfect graphs. However, using maximum flow techniques, the algorithm can be implemented to run in time $\mathcal{O}(n^4)$ (see [4] for details and further references). By Corollary 4.3, all graphs in $\mathcal{G}_2$ admit a bisimplicial elimination ordering, and so we obtain the following.

Corollary 4.5. *The* Maximum Weight Clique *problem can be solved in $\mathcal{O}(n^4)$ time for n-vertex graphs in $\mathcal{G}_2$.*

As we will show in the next section (more specifically in Theorem 5.3), this result cannot be generalized to graphs in $\mathcal{G}_k$ for $k \geq 3$, unless P = NP.

5 NP-hardness results for $\mathcal{G}_k$, $k \geq 3$

In this section we prove that for all $k \geq 3$, it is NP-hard to recognize graphs in $\mathcal{G}_k$. The time complexity of recognizing graphs in $\mathcal{G}_2$ is still unknown. We also show that the Maximum Clique problem is NP-hard for $\mathcal{G}_3$ (and consequently for $\mathcal{G}_k$ whenever $k \geq 3$).

Recall that for an integer k, the k-Coloring problem is the problem of determining whether the input graph is k-colorable. It is well known that for any integer $k \geq 3$, the k-Coloring problem is NP-complete; we prove the NP-hardness of recognizing graphs in $\mathcal{G}_k$ ($k \geq 3$) by a reduction from this problem. We begin with a technical proposition.

Proposition 5.1. *Let $k \geq 0$ be an integer, and let G be a graph. Let G' be the graph obtained from G by adding two new, nonadjacent vertices, and making them adjacent to all vertices of G. Then $\chi(\overline{G}) \leq k$ if and only if $G' \in \mathcal{G}_k$.*

Proof. Let a and b be the two vertices added to G to form G'.

Suppose first that $G' \in \mathcal{G}_k$. Clearly, $V(G)$ is a minimal (a,b)-separator of G', and so $V(G)$ is the union of k (possibly empty) cliques of G'. But then $\chi(\overline{G}) \leq k$.

Suppose now that $\chi(\overline{G}) \leq k$; we must show that $G' \in \mathcal{G}_k$. Fix two distinct, nonadjacent vertices $x, y \in V(G')$ and let S be a minimal (x,y)-separator of G'. We must show that S is the union of k (possibly empty) cliques of G'.

Suppose first that $\{x,y\} \cap \{a,b\} \neq \emptyset$. By symmetry, we may assume that $x = a$. Since b is the only nonneighbor of a in G', it follows that $y = b$. Since $\{a,b\} = \{x,y\}$ is complete to $V(G) = V(G') \setminus \{a,b\}$, it follows that $V(G)$ is the only (x,y)-separator of G'. So, $S = V(G)$. Since $\chi(\overline{G}) \leq k$, it follows that S is the union of k cliques of G'.

From now on, we assume that $\{x,y\} \cap \{a,b\} = \emptyset$, so that $x, y \in V(G)$. Note that $\{a,b\}$ is complete to $\{x,y\}$, and so $a, b \in S$. Now, $\chi(\overline{G}) \leq k$, and so $S \setminus \{a,b\}$ is the union of k cliques of G, say $C_1, \dots, C_k$. Using the fact that $\{a,b\}$ is complete to $V(G)$ in G', and the fact that $k \geq 2$, we see that S is the union of k cliques of G', namely $C_1 \cup \{a\}, C_2 \cup \{b\}, C_3, \dots, C_k$.

We have now shown that $G' \in \mathcal{G}_k$, and we are done. $\square$

Theorem 5.2. *For every integer $k \geq 3$, it is NP-hard to recognize graphs in $\mathcal{G}_k$.*

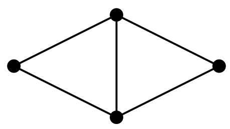

Figure 2: The diamond.

Proof. Fix an integer $k \geq 3$, and let G be any graph. We form a graph G' by adding two new, nonadjacent vertices to $\overline{G}$, and making them adjacent to all vertices of $\overline{G}$. Since $\overline{\overline{G}} = G$, Proposition 5.1 guarantees that $\chi(G) \leq k$ if and only if $G' \in \mathcal{G}_k$. Since k-Coloring is NP-complete, it follows that recognizing graphs in $\mathcal{G}_k$ is NP-hard. □

Theorem 5.3. *The* Maximum Clique *problem is NP-hard for graphs in* $\mathcal{G}_3$*.*

Proof. Note that $\mathcal{G}_3$ contains all graphs whose vertex set can be partitioned into three (possibly empty) cliques; moreover, note that the vertex set of a graph can be partitioned into three cliques if and only if the complement of the graph is 3-colorable. Thus, it suffices to show that the Maximum Stable Set problem is NP-hard for 3-colorable graphs. But this readily follows from [22]. Indeed, as observed by Poljak [22], for any graph G, the graph G^* obtained from G by subdividing each edge twice has the property that $\alpha(G^*) = \alpha(G) + |E(G)|$. But notice that for any graph G, the graph G^* is 3-colorable. Thus, since the Maximum Stable Set problem is NP-hard for general graphs, it is NP-hard for 3-colorable graphs. □

6 Diamond-free graphs in $\mathcal{G}_2$

We remind the reader that the *diamond* is the four-vertex graph obtained from the complete graph K_4 by deleting one edge (see Fig. 2). In this section, we prove a decomposition theorem for diamond-free graphs in $\mathcal{G}_2$, which implies a polynomial-time recognition algorithm for this class of graphs. We begin with some definitions.

A *hole* in a graph G is an induced cycle of G of length at least four. A *wheel* is a graph that consists of a hole and an additional vertex that has at least three neighbors in the hole. A *broken wheel* (see Fig. 3) is a wheel that consists of a hole H and an additional vertex v such that v has at least three neighbors in H, and furthermore, the neighbors of v in $V(H)$ induce a disconnected subgraph of H.

A *prism* is any subdivision of $\overline{C_6}$ (where $\overline{C_6}$ is the complement of C_6) in which the two triangles remain unsubdivided; in particular, $\overline{C_6}$ is a prism. A *pyramid* is any subdivision of the complete graph K_4 in which one triangle remains unsubdivided, and of the remaining three edges, at least two edges are subdivided at least once. A *theta* is any subdivision of the complete bipartite graph $K_{2,3}$; in particular, $K_{2,3}$ is a theta. A *3-path-configuration* (or *3PC* for short) is any theta, pyramid, or prism. The three types of 3PC are represented in Fig. 4.

The *short prism* is the graph $\overline{C_6}$, and a *long prism* is any prism other than $\overline{C_6}$. (Thus, a long prism is any prism on at least seven vertices.)

For an integer $n \geq 3$, a *short n-prism* is a graph whose vertex set can be partitioned into two n-vertex cliques, say $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$, such that for all $i, j \in \{1, \dots, n\}$, a_i is adjacent to b_j if and only if $i = j$. Note that $\overline{C_6}$ is a short 3-prism. A *complete prism* is any graph that is a short n-prism for some integer $n \geq 3$.

We will need the following decomposition theorem for (3PC, wheel)-free graphs.

Theorem 6.1 (Conforti et al. [12])**.** *If a graph G is (3PC, wheel)-free, then either G is a complete graph or a cycle, or G admits a clique cutset.*

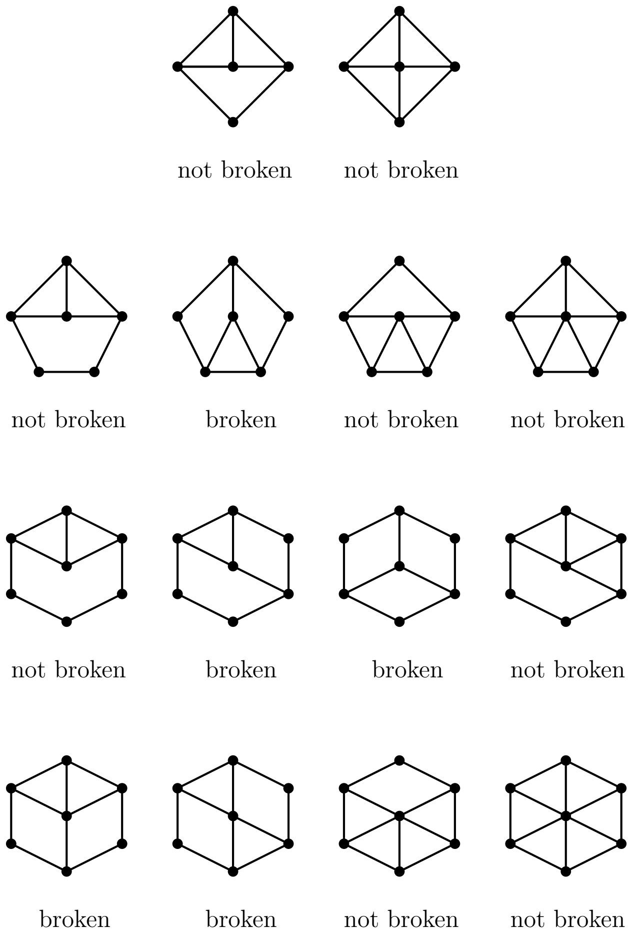

Figure 3: Some small wheels, classified as broken or not broken.

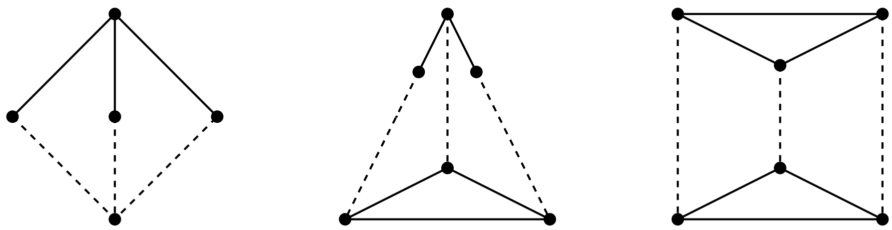

Figure 4: Three-path-configurations: theta (left), pyramid (center), and prism (right). A full line represents an edge, and a dashed line represents a path that has at least one edge.

Lemma 6.2. *Every graph in $\mathcal{G}_2$ is (theta, pyramid, long prism, broken wheel)-free.*

Proof. Since $\mathcal{G}_2$ is hereditary (by Corollary 2.2), it suffices to show that $\mathcal{G}_2$ contains no theta, no pyramid, no long prism, and no broken wheel. By Corollary 3.4, $\mathcal{G}_2$ is a subclass of the class of $K_{2,3}$-induced-minor-free graphs. Thus, it suffices to show that the class of $K_{2,3}$-induced-minor-free graphs contains no theta, no pyramid, no long prism, and no broken wheel, or equivalently, that $K_{2,3}$ is an induced minor of every theta, pyramid, long prism, or broken wheel.

Claim 1. $K_{2,3}$ is an induced minor of every theta.

Proof of Claim 1. Let H be a theta. Let a and b be distinct, nonadjacent vertices of H, and let P^1, P^2, P^3 be distinct induced paths in H, each between a and b, such that any two of P^1, P^2, P^3 have exactly two vertices (namely a and b) in common. Contracting in H all but two edges of each path P^i results in a graph isomorphic to $K_{2,3}$. This proves Claim 1. ♦

Claim 2. $K_{2,3}$ is an induced minor of every pyramid.

Proof of Claim 2. Let H be a pyramid. Let a be a vertex of H, let $B = \{b_1, b_2, b_3\}$ be a 3-vertex clique in $H \setminus a$, and let P^1, P^2, and P^3 be induced paths in H such that

- for each $i \in \{1, 2, 3\}$, the endpoints of P^i are a and b_i;
- any two of the paths P^1, P^2, P^3 have exactly one vertex (namely a) in common.

Since H is a pyramid, we know that at least two of P^1, P^2, P^3 have more than one edge; by symmetry, we may assume that P^1 and P^2 each have at least two edges. Contracting in H all but two edges of each of the paths P^1 and P^2, all but one edge of the path P^3, and the edge b_1b_2 results in a graph isomorphic to $K_{2,3}$. This proves Claim 2. ♦

Claim 3. $K_{2,3}$ is an induced minor of every long prism.

Proof of Claim 3. Let H be a long prism. Let $A = \{a_1, a_2, a_3\}$ and $B = \{b_1, b_2, b_3\}$ be disjoint 3-vertex cliques in H, and let P^1, P^2, and P^3 be induced paths in H such that

- for each $i \in \{1, 2, 3\}$, the endpoints of P^i are a_i and b_i;
- no two of the paths P^1, P^2, P^3 have any vertices in common.

Since H is a long prism, we know that at least one of P^1, P^2, P^3 has more than one edge; by symmetry, we may assume that P^1 has more than one edge. Contracting in H all but two edges of the path P^1, all but one edges of each of the paths P^2 ad P^3, and the edges a_1a_2 and b_1b_3 results in a graph isomorphic to $K_{2,3}$. This proves Claim 3. ♦

Claim 4. $K_{2,3}$ is an induced minor of every broken wheel.

Proof of Claim 4. Let W be a broken wheel, consisting of a hole $H = h_0, h_1, \dots, h_{k-1}, h_0$ (with $k \geq 4$, and indices in $\mathbb{Z}_k$) and an additional vertex v that has at least three neighbors in H, and such that the neighbors of v induce a disconnected subgraph of H. By symmetry, we may assume that v is nonadjacent to h_0 and adjacent to h_1. Let the path $h_1, \dots, h_i$ be one component of $H[N_W(v)]$. Contracting in W all edges of the hole H except for the four edges incident with h_0 and h_{i+1} results in a graph isomorphic to $K_{2,3}$. This proves Claim 4. ♦

Claims 1–4 together complete the proof of the lemma. □

Lemma 6.3. *Let G be a diamond-free graph that belongs to $\mathcal{G}_2$. Then G is (theta, pyramid, long prism, wheel)-free.*

Proof. Clearly, every wheel either contains an induced diamond or is a broken wheel. The result now follows from Lemma 6.2. □

Lemma 6.4. *Let G be a (diamond, theta, pyramid, long prism, wheel)-free graph that contains an induced $\overline{C_6}$. Then either G is a complete prism, or G admits a clique cutset.*

Proof. Recall that $\overline{C_6}$ is an induced short 3-prism; fix a maximum integer $n \geq 3$ such that G contains a short n-prism H. Set $V(H) = A \cup B$, where $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ are disjoint, n-vertex cliques, such that for all $i, j \in \{1, \dots, n\}$, a_i is adjacent to b_j in H if and only if $i = j$. We may assume that $V(H) \subsetneqq V(G)$, for otherwise, G is a complete prism, and we are done. We may further assume that G is connected, for otherwise, $\emptyset$ is a clique cutset of G, and again we are done.

Claim 1. For all $v \in V(G) \setminus V(H)$, either $N_G(v) \cap V(H) = A$, or $N_G(v) \cap V(H) = B$, or there exists some $i \in \{1, \dots, n\}$ such that $N_G(v) \cap V(H) \subseteq \{a_i, b_i\}$.

Proof of Claim 1. Fix $v \in V(G) \setminus V(H)$. We may assume that $|N_G(v) \cap V(H)| \geq 2$, for otherwise, the result is immediate. Next, if there exist distinct $i, j \in \{1, \dots, n\}$ such that v is complete to $\{a_i, b_j\}$, then $G[v, a_i, a_j, b_i, b_j]$ is either a theta or a wheel, contrary to the fact that G is (theta, wheel)-free. So, if v has a neighbor both in A and in B, then there exists some $i \in \{1, \dots, n\}$ such that $N_G(v) \cap V(H) = \{a_i, b_i\}$, and we are done. From now on, we assume that either $N_G(v) \cap V(H) \subseteq A$ or $N_G(v) \cap V(H) \subseteq B$; by symmetry, we may assume that $N_G(v) \cap V(H) \subseteq A$, and we deduce that $|N_G(v) \cap A| \geq 2$. Then v is complete to A, for otherwise, we fix pairwise distinct $a_i, a_j, a_k \in A$ such that v is adjacent to a_i, a_j and nonadjacent to a_k, and we observe that $G[v, a_i, a_j, a_k]$ is a diamond, contrary to the fact that G is diamond-free. It now follows that $N_G(v) \cap V(H) = A$, and we are done. This proves Claim 1. ♦

Claim 2. If there exists some $v \in V(G) \setminus V(H)$ such that $N_G(v) \cap V(H) = A$ (resp. such that $N_G(v) \cap V(H) = B$), then A (resp. B) is a clique cutset of G.

Proof of Claim 2. By symmetry, we may assume that some $v \in V(G) \setminus V(H)$ satisfies $N_G(v) \cap V(H) = A$; we must show that A is a clique cutset of G. By construction, A is a clique of G, and so it suffices to show that A is a cutset of G separating v from B. Suppose otherwise. Then there exists an induced path P in $G \setminus V(H)$ between v and some vertex that has a neighbor in B. Let $Q = q_0, \dots, q_t$ (with $t \geq 0$) be minimum-length subpath of P such that q_0 is complete to A and q_t has a neighbor in B; by Claim 1, $q_0 \neq q_t$, i.e., $t \geq 1$. By symmetry, we may assume that q_t is adjacent to b_1. By Claim 1, we have that either $N_G(q_t) \cap V(H) = B$, or $N_G(q_t) \cap V(H) = \{a_1, b_1\}$, or $N_G(q_t) \cap V(H) = \{b_1\}$.

Assume first that $N_G(q_t) \cap V(H) = B$. If $t = 1$, then $G[A \cup B \cup \{q_0, q_1\}]$ is a short $(n+1)$-prism, contrary to the maximality of n. So, $t \geq 2$. By the minimality of Q, all internal vertices of Q are anticomplete to B. If the internal vertices of Q are also anticomplete to A, then $G[\{a_1, a_2, b_1, b_2\} \cup V(Q)]$ is a long prism, contrary to the fact that G is long-prism-free. Hence, some internal vertex of Q has a neighbor in A; let $i \in \{1, \dots, t-1\}$ be maximum with the property that q_i has a neighbor in A. By the minimality of Q, and by Claim 1, we know that q_i has a unique neighbor in A; fix $j \in \{1, \dots, n\}$ such that a_j is the unique neighbor of q_i in A, and fix any $k \in \{1, \dots, n\} \setminus \{j\}$. But now $G[a_j, a_k, b_j, b_k, q_i, q_{i+1}, \dots, q_t]$ is a pyramid, contrary to the fact that G is pyramid-free.

Assume next that $N_G(q_t) \cap V(H) = \{a_1, b_1\}$. Let q_i be the vertex of Q with highest index such that q_i is adjacent to a_2. Then $q_i, \dots, q_t, b_1, b_2, a_2, q_i$ is a hole, and a_1 has at least three neighbors (namely, a_2, q_t, b_1) in it, contrary to the fact that G is wheel-free.

Assume finally that $N_G(q_t) \cap V(H) = \{b_1\}$. Let q_i (resp. q_j) be the vertex of Q with highest index such that q_i (resp. q_j) is adjacent to a_2 (resp. a_1). If $j \geq i$ then $q_i, \dots, q_t, b_1, b_2, a_2, q_i$ is a hole and a_1 has at least three neighbors in it (namely, a_2, q_j, b_1), contrary to the fact that G is wheel-free. So, $i > j$. Then $q_j, \dots, q_t, b_1, a_1, q_j$ is a hole and a_2 has two nonadjacent neighbors in it (namely, a_1, q_i), and hence $G[q_j, \dots, , q_t, b_1, a_1, a_2]$ is a theta or a wheel, contrary to the fact that G is (theta, wheel)-free. This proves Claim 2. ♦

In view of Claims 1 and 2, we assume from now on that for all $v \in V(G) \setminus V(H)$, there exists some $i \in \{1, \dots, n\}$ such that $N_G(v) \cap V(H) \subseteq \{a_i, b_i\}$.

Claim 3. For all $i \in \{1, \dots, n\}$, if some vertex in $V(G) \setminus V(H)$ is complete to $\{a_i, b_i\}$, then $\{a_i, b_i\}$ is a clique cutset of G.

Proof of Claim 3. By symmetry, we may assume that some vertex of $V(G) \setminus V(H)$ is complete to $\{a_1, b_1\}$; we must show that $\{a_1, b_1\}$ is a clique cutset of G. Since a_1 is adjacent to b_1, it suffices to show that $G \setminus \{a_1, b_1\}$ is disconnected. Suppose otherwise. Then there exists an induced path $P = p_0, \dots, p_s$ (with $s \geq 0$) in $G \setminus (A \cup B)$ such that p_0 is complete to $\{a_1, b_1\}$ and p_s has a neighbor in $V(H) \setminus \{a_1, b_1\}$; we may assume that the path P was chosen so that its length is minimum. By Claim 1, we have that $N_G(p_0) \cap V(H) = \{a_1, b_1\}$, and so $s \geq 1$. Furthermore, by the minimality of P, $\{p_0, \dots, p_{s-1}\}$ is anticomplete to $V(H) \setminus \{a_1, b_1\}$. By symmetry, we may assume that $a_2 \in N_G(p_s) \cap V(H) \subseteq \{a_2, b_2\}$. Let i be the largest index in $\{0, \dots, s\}$ such that p_i is adjacent to b_1 (such an i exists because p_0 is adjacent to b_1). Now $p_i, \dots, p_s, a_2, a_3, b_3, b_1, p_i$ is a hole, and a_1 has at least three neighbors (namely, a_2, a_3, b_1) in it, contrary to the fact that G is wheel-free. This proves Claim 3. ♦

In view of Claim 3, we may now assume that no vertex in $V(G) \setminus V(H)$ has more than one neighbor in $V(H)$. Let N_A be the set of all vertices in $V(G) \setminus V(H)$ that have a neighbor in A, and let N_B be the set of all vertices in $V(G) \setminus V(H)$ that have a neighbor in B. Then $N_A \cap N_B = \emptyset$. Since $V(H) \subsetneqq V(G)$ and G is connected, we have that $N_A \cup N_B \neq \emptyset$. If $N_A = \emptyset$, then B is a clique cutset of G, and if $N_B = \emptyset$, then A is a clique cutset of G. So, we may assume that N_A and N_B are both nonempty. Furthermore, we may assume that there is an induced path in $G \setminus V(H)$ between N_A and N_B, for otherwise, both A and B are clique cutsets of G, and we are done. Let $P = p_0, \dots, p_s$ (with $s \geq 0$) be a minimum-length path in $G \setminus V(H)$ such that $p_0 \in N_A$ and $p_s \in N_B$; since $N_A \cap N_B = \emptyset$, we see that $s \geq 1$. Furthermore, the minimality of P implies that the interior of P is anticomplete to $V(H)$.

Since no vertex of $V(G) \setminus V(H)$ has more than one neighbor in $V(H)$, we may assume by symmetry that $N_G(p_0) \cap V(H) = \{a_1\}$, and that either $N_G(p_s) \cap V(H) = \{b_1\}$ or $N_G(p_s) \cap V(H) = \{b_2\}$. But if $N_G(p_s) \cap V(H) = \{b_2\}$, then $G[a_1, a_2, b_1, b_2, p_0, \dots, p_s]$ is a theta, contrary to the fact that G is theta-free. So, $N_G(p_s) \cap V(H) = \{b_1\}$. We now have that $V(P)$ is anticomplete to $V(H) \setminus \{a_1, b_1\}$.

Our goal is to show that $\{a_1, b_1\}$ is a clique cutset of G. Suppose otherwise; then $G \setminus \{a_1, b_1\}$ is connected. Since $V(P)$ is anticomplete to $V(H) \setminus \{a_1, b_1\}$, we see that there exists an induced path $Q = q_0, \dots, q_t$ (with $t \geq 0$) in $G \setminus (V(H) \cup V(P))$ such that q_0 has a neighbor in $V(P)$, and q_t has a neighbor in $V(H) \setminus \{a_1, b_1\}$; we may assume that Q is a minimum-length path with this property, so that q_0 is the only vertex of Q with a neighbor in $V(P)$, and q_t is the only vertex of Q with a neighbor in $V(H) \setminus \{a_1, b_1\}$. By symmetry, we may further assume that q_t is adjacent to a_2; then $N_G(q_t) \cap V(H) = \{a_2\}$. Let i be the largest index in $\{0, \dots, s\}$ such that q_0 is adjacent to p_i. Then $p_i, \dots, p_s, b_1, b_3, a_3, a_2, q_t, \dots, q_0, p_i$ is a hole in G, and b_2 has at least three neighbors (namely, a_2, b_1, b_3) in it, contrary to the fact that G is wheel-free. This completes the proof. □

Lemma 6.5. *Let G be a (diamond, theta, pyramid, long prism, wheel)-free graph. Then either G is a complete prism, a hole, or a complete graph, or G admits a clique cutset.*

Proof. If G contains an induced $\overline{C_6}$, then the result follows from Lemma 6.4. Otherwise, we have that G is (3PC, wheel)-free, and the result follows from Theorem 6.1. □

Theorem 6.6. *Let G be a diamond-free graph that belongs to $\mathcal{G}_2$. Then either G is a complete prism, a cycle, or a complete graph, or G admits a clique cutset.*

Proof. This follows immediately from Lemmas 6.3 and 6.5. □

6.1 Algorithmic considerations

Clearly, complete prisms, cycles, and complete graphs are diamond-free and belong to $\mathcal{G}_2$, and furthermore, they can all be recognized in polynomial time. So, Corollary 2.4 and Theorem 6.6, together with Tarjan's [24] tool for handling clique cutsets, allow us to recognize diamond-free graphs in $\mathcal{G}_2$ in polynomial time.

Proposition 6.7. *There exists an algorithm running in time $\mathcal{O}(n(n+m))$ that correctly determines if an input graph G with n vertices and m edges is a diamond-free graph in $\mathcal{G}_2$.*

Proof. Given a graph G with n vertices and m edges, testing if G is diamond-free can be done in time $\mathcal{O}(n(n+m))$ by verifying if for each vertex v its neighborhood $N_G(v)$ induces a disjoint union of complete graphs. Assuming G is diamond-free, we compute the connected components of G and run Tarjan's algorithm [24] on each nontrivial component of G. This can again be done in time $\mathcal{O}(n(n+m))$. Tarjan's algorithm produces a family $\mathcal{H}$ of $\mathcal{O}(n)$ induced subgraphs of G that do not have any clique cutsets. We then check, for each graph $H \in \mathcal{H}$, whether H is a complete prism, cycle, or a complete graph. If this is the case, the algorithm determines that G belongs to $\mathcal{G}_2$, and otherwise, it determines that G does not belong to $\mathcal{G}_2$. The correctness follows from Corollary 2.4 and Theorem 6.6.

To complete the proof, we show that testing whether a given $H \in \mathcal{H}$ satisfies one of the desired properties can be done in time $\mathcal{O}(n+m)$. Since H is connected, testing if it is a cycle or a complete graph can be done in linear time simply by checking if all the vertex degrees are equal to 2 or to $|V(H)|-1$, respectively. If none of these cases occurs, we can assume that $n = 2k$ for some $k \geq 3$ and that every vertex in H has degree exactly k, since otherwise, we can infer that H is not a complete prism. We choose an arbitrary vertex $v \in V(H)$ and compute the components of the graph $H[N_H(v)]$. If H is a short k-prism, then $H[N_H(v)]$ has exactly two components, say C and D, such that C is isomorphic to a complete graph K_{k-1} and D is a trivial component. Repeating the computation of components for the subgraph of H induced by the neighborhood of the unique vertex in D, we identify the two cliques A and B, each of size k, that partition $V(H)$ (more precisely, $A = C \cup \{v\}$ and B is obtained analogously.) It remains to verify if each vertex in A has a unique neighbor in B and vice versa. If this is the case, then H is a complete prism, otherwise it is not. Each of the above constantly many steps can be carried out in linear time. □

Moreover, it is clear that the Maximum Weight Clique, Maximum Weight Stable Set, and Vertex Coloring can be solved in polynomial time for complete prisms, cycles, and complete graphs. Thus, Theorem 6.6 and the algorithm from [24] allow us to solve these three optimization problems in polynomial time for diamond-free graphs in $\mathcal{G}_2$. A more precise time complexity analysis is provided by the following.

Theorem 6.8. *When restricted to the class of diamond-free graphs in $\mathcal{G}_2$ with n vertices and m edges, the* Maximum Weight Clique *and* Vertex Coloring *problems can be solved in $\mathcal{O}(n(n+m))$ time and the* Maximum Weight Stable Set *problem in $\mathcal{O}(n^2(n+m))$ time.*

Proof. Let G be a diamond-free graph with n vertices and m edges that belongs to $\mathcal{G}_2$. For the Maximum Weight Clique and Vertex Coloring problems, the approach is as follows. Using Tarjan's algorithm [24], we compute in time $\mathcal{O}(n(n+m))$ a family $\mathcal{H}$ of $\mathcal{O}(n)$ induced subgraphs of G that do not have any clique cutsets. We now iterate over all $H \in \mathcal{H}$ and solve the Maximum Weight Clique and Vertex Coloring problems on H in linear time. By Theorem 6.6, each graph $H \in \mathcal{H}$ is a complete prism, cycle, or a complete graph; as explained in the proof of Proposition 6.7, which of these cases occurs can be determined in linear time in the size of H. If H is a complete graph, then its chromatic number is $|V(H)|$ and its vertex set solves the Maximum Weight Clique problem. Otherwise, if H is a cycle with at least four vertices, then its chromatic number is either 2 or 3, depending on whether $|V(H)|$ is even or odd, respectively, and to solve the Maximum Weight Clique problem, we only need to examine its edges. Finally, if H is a complete prism, say with $|V(H)| = 2k$ for some $k \geq 3$, then we can identify in linear time the two cliques A and B, each of size k, that partition $V(H)$. Then, the chromatic number of H is k, and to solve the Maximum Weight Clique problem, we only need to examine the two cliques A and B and the k edges connecting them. In all these cases, the Maximum Weight Clique and Vertex Coloring problems can be solved in linear time for graphs in $\mathcal{H}$. Since each clique of G is fully contained in one of the graphs in $\mathcal{H}$, this provides an efficient solution to the Maximum Weight Clique problem on G. Similarly, the chromatic number of G is the maximum chromatic number of the graphs in $\mathcal{H}$.

For the Maximum Weight Stable Set problem, the approach is similar, except that in each decomposition step decomposes G along a cut partition (A, B, C) of G such that C is a clique and the subgraph of G induced by $A \cup C$ belongs to $\mathcal{H}$. For each vertex $v \in C$, we determine a maximum-weight stable set of the subgraph of G induced by $A \setminus N_G(v)$, a maximum-weight stable set of the subgraph of G induced by A, redefine the weights on C, and solve the problem recursively on the graph $G - A$. (We refer to [24] for details; see also [8, Section 8.1].) Thus, we solve $\mathcal{O}(|V(G)|)$ subproblems for each graph $H \in \mathcal{H}$ for a total of $\mathcal{O}(|V(G)|^2)$ subproblems. Each of these subproblems can be solved in linear time. If H is a complete graph, then a heaviest vertex forms a maximum-weight stable set. If H is a cycle, then we can use the fact that cycles have bounded treewidth and apply the results from [2, 6]. Finally, if H is a complete prism, then H has clique-width at most 4, and a 4-expression of H can be obtained in linear time (see [10]) and hence the result of Courcelle, Makowsky, and Rotics [13] applies. □

In conclusion, let us put the results of Theorem 6.8 in perspective by comparing them with the known complexities of the three problems in the larger classes of diamond-free graphs and graphs in $\mathcal{G}_2$. First, the Vertex Coloring problem is NP-hard for diamond-free graphs [20], as well as for graphs in $\mathcal{G}_2$, since it is already hard for the subclass of circular-arc graphs [17]. The situation is somewhat different for the Maximum Weight Stable Set problem, which is NP-hard (even in the unweighted case) in the class of diamond-free graphs, as can be seen using Poljak's reduction [22], but solvable in $\mathcal{O}(n^6)$ time for n-vertex graphs in $\mathcal{G}_2$ (see Corollary 3.5). Finally, while the Maximum Weight Clique problem is known to be solvable in polynomial time both for diamond-free graphs as well as for graphs in $\mathcal{G}_2$, the running time of the algorithm given by Theorem 6.8 improves on both time complexities. By Corollary 4.5, the problem can be solved in $\mathcal{O}(n^4)$ time for n-vertex graphs in $\mathcal{G}_2$. For the class of diamond-free graphs, observe that every edge in such a graph is contained in a unique maximal clique. Thus, a diamond-free graph with n vertices and m edges has $\mathcal{O}(n+m)$ maximal cliques, and the Maximum Weight Clique problem can be solved in polynomial time by enumerating all maximal cliques and returning one of maximum weight. Using, for example, the maximal clique enumeration algorithm due to Makino and Uno [21], this would result in an overall running time of $\mathcal{O}(n^{2.37}(n+m))$ on diamond-free graphs with n vertices and m edges.

References

[1] P. Aboulker, P. Charbit, N. Trotignon, and K. Vušković. Vertex elimination orderings for hereditary graph classes. *Discrete Mathematics*, 338(5):825–834, 2015.

[2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308–340, 1991.

[3] J. Backer. Separator orders in interval, cocomparability, and AT-free graphs. *Discrete Applied Mathematics*, 159(8):717–726, 2011.

[4] J. Beisegel, M. Chudnovsky, V. Gurvich, M. Milanič, and M. Servatius. Avoidable vertices and edges in graphs: existence, characterization, and applications. *Discrete Appl. Math.*, 309:285–300, 2022.

[5] A. Berry and J.-P. Bordat. Separability generalizes Dirac's theorem. *Discrete Applied Mathematics*, 84(1-3):43–53, 1998.

[6] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.

[7] H. L. Bodlaender and A. M. C. A. Koster. Safe separators for treewidth. *Discrete mathematics*, 306(3):337–350, 2006.

[8] V. Boncompagni, I. Penev, and K. Vušković. Clique-cutsets beyond chordal graphs. *J. Graph Theory*, 91(2):192–246, 2019.

[9] V. Bouchitté and I. Todinca. Treewidth and minimum fill-in: grouping the minimal separators. *SIAM Journal on Computing*, 31(1):212–232, 2001.

[10] A. Brandstädt, F. F. Dragan, H.-O. Le, and R. Mosca. New graph classes of bounded clique-width. *Theory Comput. Syst.*, 38(5):623–645, 2005.

[11] S. Chaplick, M. Töpfer, J. Voborník, and P. Zeman. On H-topological intersection graphs. *Algorithmica*, 83(11):3281–3318, 2021.

[12] M. Conforti, G. Cornuéjols, A. Kapoor, and K. Vušković. Universally signable graphs. *Combinatorica*, 17(1):67–77, 1997.

[13] B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.

[14] C. Dallard, M. Milanič, and K. Štorgel. Tree decompositions with bounded independence number and their algorithmic applications. *arXiv:2111.04543*, 2021.

[15] C. Dallard, M. Milanič, and K. Štorgel. Treewidth versus clique number. I. Graph classes with a forbidden structure. *SIAM J. Discrete Math.*, 35(4):2618–2646, 2021.

[16] G. A. Dirac. On rigid circuit graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 25:71–76, 1961.

[17] M. Garey, D. Johnson, G. Miller, and C. Papadimitriou. The complexity of coloring circular arcs and chords. *SIAM Journal on Algebraic Discrete Methods*, 1(2):216–227, 1980.

[18] T. R. Hartinger and M. Milanič. Partial characterizations of 1-perfectly orientable graphs. *J. Graph Theory*, 85(2):378–394, 2017.

[19] P. Heggernes. Minimal triangulations of graphs: a survey. *Discrete Mathematics*, 306(3):297–317, 2006.

[20] D. Kráľ, J. Kratochvíl, Z. Tuza, and G. J. Woeginger. Complexity of coloring graphs without forbidden induced subgraphs. In *Graph-theoretic concepts in computer science (Boltenhagen, 2001)*, volume 2204 of *Lecture Notes in Comput. Sci.*, pages 254–262. Springer, Berlin, 2001.

[21] K. Makino and T. Uno. New algorithms for enumerating all maximal cliques. In *Algorithm theory—SWAT 2004*, volume 3111 of *Lecture Notes in Comput. Sci.*, pages 260–272. Springer, Berlin, 2004.

[22] S. Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15(2):307–309, 1974.

[23] D. J. Rose, R. E. Tarjan, and G. S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM Journal on computing*, 5(2):266–283, 1976.

[24] R. E. Tarjan. Decomposition by clique separators. *Discrete mathematics*, 55(2):221–232, 1985.

[25] Y. Ye and A. Borodin. Elimination graphs. *ACM Transactions on Algorithms*, 8(2):Article 14, 1–23, 2012.