AIMS Mathematics, 11(11): 37133–37157.
DOI:10.3934/math.20261435
Received: 31 August 2026
Revised: 24 October 2026
Accepted: 02 November 2026
Published: 10 November 2026

https://www.aimspress.com/journal/Math

Research article

k-Independence in P_4-tidy graphs

Márcia R. Cappelle* **and Luan César D. Carvalho**

Instituto de Informática, Universidade Federal de Goiás, Goiânia-GO, Brazil

* **Correspondence:** E-mail: mcappelle@ufg.br.

Abstract: A set $S \subseteq V(G)$ of a graph G is k-independent if every vertex of S has at most $k-1$ neighbors in S, where $k \in \mathbb{Z}^+$. The k-independence number $\alpha_k(G)$ is the maximum cardinality of a k-independent set of G; determining it is $\mathcal{NP}$-hard for general graphs. We study this parameter on P_4-tidy graphs, a class that properly contains cographs and P_4-sparse graphs. Exploiting the structural decomposition of P_4-tidy graphs, we derive closed formulas for α_k under disjoint union and join, as well as for all types of spiders and quasi-spiders, expressed in terms of the k-independence numbers of the graphs arising in the decomposition. These formulas yield a dynamic programming algorithm over the decomposition tree that computes $\alpha_k(G)$ in $O(k^3n + |E(G)|)$ time, with k part of the input. Although polynomial-time solvability on this class already follows from general algorithms for graphs of bounded clique-width, those algorithms run in $n^{O(w)}$ time for clique-width w and provide no explicit formulas. Our results extend the closed formulas and the efficient algorithm previously known for P_4-sparse graphs to the full class of P_4-tidy graphs.

Keywords: P_4-tidy graphs; spiders; decomposition tree; polynomial-time algorithm; bounded-degree vertex deletion
Mathematics Subject Classification: 05C69, 05C85

1. Introduction

A set $S \subseteq V(G)$ is called k-independent if the maximum degree of the subgraph induced by S is at most $k-1$. The maximum cardinality of a k-independent set of G is the k-independence number, denoted by $\alpha_k(G)$. Thus, the classical independence number corresponds to the case $k = 1$. Equivalently, $\alpha_k(G)$ is the maximum order of an induced subgraph of G with maximum degree at most $k-1$. We adopt the convention that $\alpha_k(\emptyset) = 0$ for every $k \geq 1$. This generalization of independence was introduced by Fink and Jacobson [1], and determining $\alpha_k(G)$ is $\mathcal{NP}$-hard for general graphs [2]. Since then, k-independence has been investigated from both structural and extremal perspectives, including bounds for $\alpha_k(G)$, relations with other graph parameters, and results for specific graph classes [3–5].

On the complexity side, the problem remains $\mathcal{NP}$-hard even on restricted inputs such as Cartesian products and complementary prisms [6].

Graph classes defined by restrictions on induced P_4's provide a natural setting for this question. At the most restrictive level, cographs are exactly the P_4-free graphs and admit a recursive decomposition by disjoint union and join. This decomposition is responsible for the efficient solution of several optimization problems on cographs, including the maximum independent set problem. Several extensions of cographs have been introduced by allowing a controlled number or arrangement of induced P_4's. A graph is P_4-sparse [7] if no set of five vertices induces more than one P_4 and P_4-tidy [8] if, for every vertex set A inducing a P_4, there is at most one vertex $v \notin A$ such that $A \cup \{v\}$ induces at least two P_4's. Cographs form a proper subclass of P_4-sparse graphs, which in turn form a proper subclass of P_4-tidy graphs.

The class of P_4-tidy graphs is particularly suitable for algorithmic investigation because it admits a strong structural characterization. Roughly speaking, every P_4-tidy graph can be recursively decomposed by disjoint union and join or has the structure of a spider or quasi-spider, apart from a small number of exceptional graphs. Moreover, P_4-tidy graphs can be recognized in linear time [8]. This decomposition suggests a dynamic programming approach in which the k-independence number of a graph is obtained from the corresponding values of its components or modules. However, unlike ordinary independence, for $k \geq 2$ the behavior of a k-independent set under the join operation and within spiders and quasi-spiders depends on the internal degrees of the selected vertices, making the extension nontrivial.

The decomposition has been exploited both to obtain efficient algorithms for domination-type problems on P_4-tidy graphs [9, 10] and to characterize structural properties of the class, such as well-coveredness [11] and the existence of open-independent, open-locating-dominating sets [12].

The k-independence number of P_4-sparse graphs was previously studied by the authors in [13]. That work established the formula for the join, closed formulas for thin and thick spiders, and a dynamic programming algorithm over the P_4-sparse decomposition tree. These results form the starting point of the present paper, where the join formula receives a complete proof (Theorem 7, including the attainability of every term) and the spider formulas are derived from a single exchange principle (Lemma 9). The results of Section 3 (Observations 3 and 4, Lemma 5, Corollaries 6 and 8 and the formula of Theorem 7) and Theorem 11 already appear in [13], possibly in a slightly weaker form or with a different proof, and Lemma 10 extends the normalization for spiders given there. All other results, from Lemma 9 to Theorem 22, are new. The extension to the full class of P_4-tidy graphs requires the treatment of additional structures that do not occur in the P_4-sparse decomposition. In particular, we derive new normalization lemmas and closed formulas for the eight types of quasi-spiders, according to the location of the duplicated vertex, the true/false-twin relation, and the thin/thick type of the underlying spider. We also incorporate the exceptional graphs P_5, $\overline{P_5}$, and C_5. These new structural results are then integrated into the dynamic programming framework to obtain an algorithm for the full class of P_4-tidy graphs. Since P_4-tidy graphs properly contain cographs and P_4-sparse graphs, this extends the explicit formulas and the efficient algorithm to a substantially broader class of graphs with few induced P_4's.

P_4-tidy graphs have clique-width at most four, and a corresponding clique-width expression can be obtained in linear time [14]. Computing $\alpha_k(G)$ is equivalent to the Bounded-Degree Vertex Deletion problem with degree bound $k-1$: a set S is k-independent if and only if deleting $V(G) \setminus S$ leaves a

graph of maximum degree at most $k-1$. The same problem has also been studied under the name of maximum k-dependent set; in the convention most commonly used there, a set is k-dependent if every vertex has at most k neighbors in it, so the k-independent sets considered here are exactly the $(k-1)$-dependent sets, and $\alpha_k(G)$ is the maximum cardinality of a $(k-1)$-dependent set of G. Bounded-Degree Vertex Deletion is W[1]-hard when parameterized by treewidth [15]. For every fixed k, the problem is expressible in MSO_1, so the metatheorem of Courcelle, Makowsky, and Rotics [14] yields a linear-time algorithm on graphs of bounded clique-width; the formula, however, depends on k, and the hidden constant is not elementary in its size. Still for fixed k, polynomial-time solvability also follows from the vertex-partitioning algorithms of Gerber and Kobler [16]. When the degree bound is part of the input, the problem can still be solved exactly by dynamic programming over a clique-width expression with w labels, in $n^{O(w)}$ time. Lampis [17] takes such an exact dynamic program as the starting point for the bicriteria approximation scheme developed there, which, given a clique-width expression with w labels, runs in $(\log n/\varepsilon)^{O(w)}n^{O(1)}$ time and returns a set no larger than an optimal one whose removal leaves maximum degree at most $(1+\varepsilon)$ times the prescribed bound. Consequently, polynomial-time solvability of the k-independence problem on P_4-tidy graphs is already implied by these general results. Our contribution is therefore of a different nature: we derive explicit closed formulas for every structure occurring in the P_4-tidy decomposition, which yield a direct dynamic programming algorithm with running time $O(k^3n + |E(G)|)$. This gives an explicit polynomial dependence on k and a substantially sharper running-time bound than the generic dynamic program over clique-width expressions.

The remainder of the paper is organized as follows: Section 2 introduces the necessary definitions and structural properties of P_4-tidy graphs. Section 3 studies k-independence under disjoint union and join. Section 4 establishes the results for spiders and quasi-spiders, and Section 5 presents the dynamic programming algorithm and its complexity analysis. Finally, Section 6 concludes the paper and discusses directions for future research.

2. Definitions and preliminary results

Throughout this paper, $G = (V(G), E(G))$ is assumed to be a finite simple undirected graph with vertex set $V(G)$ and edge set $E(G)$. A subgraph of G with n vertices that is a path (a cycle) is termed an n-path (an n-cycle), and is denoted by P_n (C_n). The open neighborhood of a vertex $v \in V(G)$ is denoted by $N_G(v) = \{u \in V(G) \mid uv \in E(G)\}$, and its closed neighborhood is denoted by $N_G[v] = N_G(v) \cup \{v\}$. Let $S \subseteq V(G)$. The subgraph of G induced by S is denoted by $G[S]$.

The degree of a vertex v of G, denoted by $d(v)$ or $d_G(v)$, is the cardinality of this open neighborhood, i.e., $d(v) = |N_G(v)|$. If $S \subseteq V(G)$ and $v \in S$, we denote by $d_S(v)$ the degree of v in $G[S]$. A vertex v with $d(v) = 0$ is an isolated vertex. The minimum and maximum degrees of the vertices of a graph G are denoted, respectively, by $\delta(G)$ and $\Delta(G)$. A clique C of a graph G is a subset of vertices of G such that its vertices are pairwise adjacent in G. Two vertices u and v of a graph are true twins if $N[u] = N[v]$ (in particular, u and v are adjacent) and false twins if $N(u) = N(v)$ (in particular, u and v are nonadjacent). We use $n(G)$ to denote the order of G and write $n = n(G)$ when the graph under consideration is clear from the context.

Let G_1 and G_2 be two graphs such that $V(G_1) \cap V(G_2) = \emptyset$. We denote the disjoint union (resp. join) of G_1 and G_2 by $G_1 \oplus G_2$ (resp. $G_1 \otimes G_2$), and the complement graph of a graph G by $\overline{G}$. A graph G is termed co-connected if $\overline{G}$ is connected. If $X \subseteq V(G)$, the subgraph of G obtained by removing both the

vertices in X and the edges incident with them in G is denoted by $G \setminus X$. For a graph H and a natural number p, the graph obtained by the disjoint union of p copies of H is denoted by pH. For integer $\ell \geq 1$, the set $\{1, \ldots, \ell\}$ is denoted by $[\ell]$.

A cograph can be constructed recursively from isolated vertices by applying disjoint union and join operations. A consequence of this fact is that, for every cograph H with at least two vertices, H or $\overline{H}$ is disconnected.

A graph is P_4-sparse if no five vertices induce more than one P_4 [7]. The next theorem provides an alternative characterization of P_4-sparse graphs.

Theorem 1. *[18] Let G be a graph. Then, G is a P_4-sparse graph if and only if, for every induced subgraph H of G with at least two vertices, exactly one of the following conditions holds:*

1) H is disconnected;
2) $\overline{H}$ is disconnected;
3) H is isomorphic to a spider.

Recall from the introduction that a graph $G = (V, E)$ is P_4-tidy if, for every vertex set A inducing a P_4 in G, there is at most one vertex $v \in V \setminus A$ such that $G[A \cup \{v\}]$ contains at least two induced P_4's. This class includes spider and quasi-spider graphs, which we now define.

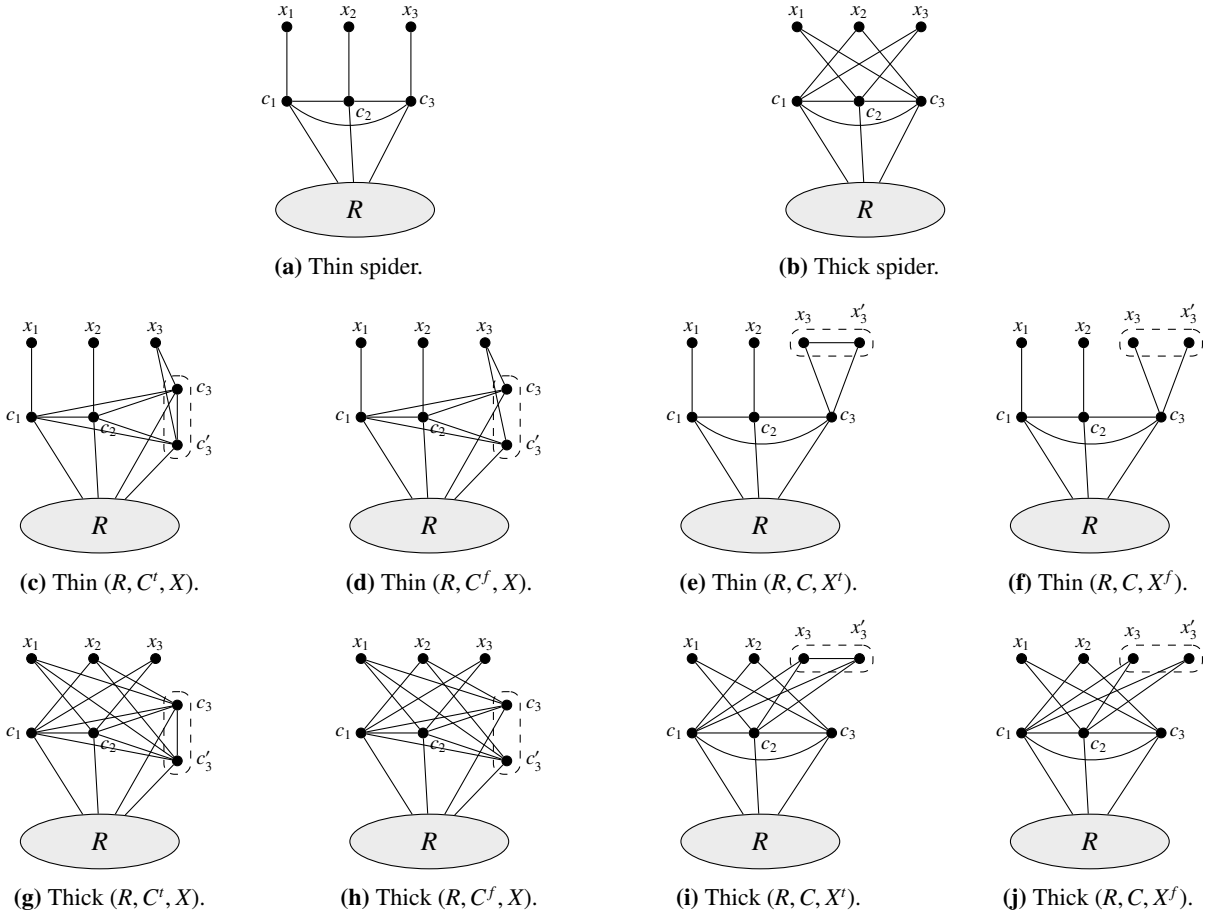
A spider is a graph whose vertex set has a partition (R, C, X), where $C = \{c_1, \ldots, c_t\}$ and $X = \{x_1, \ldots, x_t\}$, for some $t \geq 2$, are respectively a clique and an independent set; x_i is adjacent to c_j if and only if $i = j$ (a thin spider), or x_i is adjacent to c_j if and only if $i \neq j$ (a thick spider); and every vertex of R is adjacent to each vertex of C and is not adjacent to any vertex of X. The head R may be empty. The size t of both C and X is called the weight of the spider, and the set R in the partition is called its head. Notice that if $t = 2$, then the thick and thin spider graphs are isomorphic. In addition, the complement of a thin spider is a thick spider, and vice versa.

A quasi-spider is obtained from a spider by replacing exactly one vertex of $C \cup X$ by either a K_2, producing a pair of true twins, or a $\overline{K_2}$, producing a pair of false twins. According to the location and type of the duplicated vertex, we use the notation

$$(R, C^t, X), \qquad (R, C^f, X), \qquad (R, C, X^t), \qquad (R, C, X^f),$$

where the position of the superscript indicates whether the duplicated vertex belongs to C or to X. A quasi-spider is called thin or thick according to whether the underlying spider is thin or thick, respectively. Thus, each of the four types above may be either thin or thick.

The Figure 1 is a structural theorem for P_4-tidy graphs in terms of spider and quasi-spider graphs. Spiders and quasi-spiders are co-connected graphs.



(a) Thin spider. **(b)** Thick spider.

(c) Thin (R, C^t, X). **(d)** Thin (R, C^f, X). **(e)** Thin (R, C, X^t). **(f)** Thin (R, C, X^f).

(g) Thick (R, C^t, X). **(h)** Thick (R, C^f, X). **(i)** Thick (R, C, X^t). **(j)** Thick (R, C, X^f).

Figure 1. Spiders and quasi-spiders of weight 3. In the underlying spider, $C = \{c_1, c_2, c_3\}$ is a clique, $X = \{x_1, x_2, x_3\}$ is an independent set, and the head R (possibly empty) is complete to C and anticomplete to X. In each quasi-spider, the vertex c_3 or x_3 of the underlying spider is replaced by a twin pair, enclosed by a dashed box: its two vertices have the same neighbors outside the pair and are adjacent (true twins) or nonadjacent (false twins).

Theorem 2 ([8]). *G is a P_4-tidy graph if and only if exactly one of the following statements holds:*

1) G is the disjoint union or the join of two P_4-tidy graphs;
2) G is a spider or a quasi-spider graph with partition (R, C, X) such that either R induces a P_4-tidy graph or is empty;
3) G is isomorphic to C_5, P_5, $\overline{P_5}$, or K_1.

The structural characterization of P_4-tidy graphs naturally gives rise to a decomposition tree. For a P_4-tidy graph G, we denote by $T(G)$ a rooted tree in which each node represents an induced subgraph of G. The root represents G. Each node of $T(G)$ is a single vertex, an exceptional node representing one of the exceptional graphs P_5, $\overline{P_5}$, and C_5, or a node representing a disjoint union, a join, a spider, or a quasi-spider. The children of a union or join node correspond to its operands, whereas a spider or quasi-spider node with partition (R, C, X) has a single child, corresponding to the head R, or no child when $R = \emptyset$. Thus, processing $T(G)$ in bottom-up order ensures that the values associated with the proper subgraphs required at a node have already been computed. The construction of $T(G)$, together

with the labels used by the algorithm, is given in Lemma 21.

3. k-Independence in graphs obtained with disjoint union and join operations

If G is a disconnected graph, its k-independence number is the sum of the k-independence numbers of its components. Furthermore, if a graph G has at most k vertices, then $\alpha_k(G) = n(G)$. We state these immediate facts in the following observations.

Observation 3. *Let G be a graph and $\ell \geq 2$. If $G = G_1 \oplus \cdots \oplus G_\ell$, then $\alpha_k(G) = \sum_{i=1}^{\ell} \alpha_k(G_i)$.*

Observation 4. *Let G be a graph. If $n(G) \leq k$, then $\alpha_k(G) = n(G)$.*

We now consider k-independent sets in graphs obtained by the join operation. We first establish a bound for sets containing vertices from both sides of the join.

Lemma 5. *Let $k \geq 2$, let $G = G_1 \otimes G_2$, and let $S = S_1 \cup S_2$ be a k-independent set of G, where $S_i \subseteq V(G_i)$ and $S_1, S_2 \neq \emptyset$. Let $x = \Delta(G_1[S_1]) + 1$ and $y = \Delta(G_2[S_2]) + 1$. Then $x, y \in [k-1]$, $|S_1| \leq k - y$, and $|S_2| \leq k - x$.*

Proof. Since $S_2 \neq \emptyset$, every vertex of S_1 has at least one neighbor in S_2. Since S is k-independent, it follows that $\Delta(G_1[S_1]) \leq k-2$, and hence $x \in [k-1]$. Analogously, $y \in [k-1]$. Let $v \in S_1$ be such that $d_{S_1}(v) = x - 1$. Since every vertex of G_1 is adjacent to every vertex of G_2, we have $d_S(v) = x - 1 + |S_2| \leq k - 1$, and therefore $|S_2| \leq k - x$. By symmetry, $|S_1| \leq k - y$. □

As an immediate consequence, we obtain an upper bound on the cardinality of a k-independent set containing vertices from both sides of a join.

Corollary 6. *Let $k \geq 2$, let $G = G_1 \otimes G_2$, and let $S = S_1 \cup S_2$ be a k-independent set of G, where $S_i \subseteq V(G_i)$ and $S_1, S_2 \neq \emptyset$. Then $|S| \leq 2(k-1)$. Moreover, if $|S| = 2(k-1)$, then $|S_1| = |S_2| = k - 1$ and both S_1 and S_2 are independent sets.*

Proof. By the previous lemma, $|S| \leq (k-y) + (k-x) = 2k - x - y \leq 2(k-1)$, since $x, y \geq 1$. If $|S| = 2(k-1)$, then necessarily $x = y = 1$. Thus $\Delta(G_1[S_1]) = \Delta(G_2[S_2]) = 0$, and hence S_1 and S_2 are independent sets. Moreover, $|S_1| \leq k-1$ and $|S_2| \leq k-1$, so equality in the sum implies $|S_1| = |S_2| = k - 1$. □

We can now give an exact formula for the k-independence number of the join of two graphs.

Theorem 7. *Let $G = G_1 \otimes G_2$. If $k = 1$, then $\alpha_1(G) = \max\{\alpha_1(G_1), \alpha_1(G_2)\}$, and for every $k \geq 2$,*

$$\alpha_k(G) = \max\left\{\alpha_k(G_1), \alpha_k(G_2), \max_{x,y \in [k-1]}\left\{\min\{\alpha_x(G_1), k-y\} + \min\{\alpha_y(G_2), k-x\}\right\}\right\}.$$

Proof. For $k = 1$, an independent set of a join cannot contain vertices from both sides, and hence $\alpha_1(G) = \max\{\alpha_1(G_1), \alpha_1(G_2)\}$.

Now, suppose that $k \geq 2$. Let S be a maximum k-independent set of G, and let $S_i = S \cap V(G_i)$ for $i \in \{1, 2\}$. If $S_2 = \emptyset$, then $|S| \leq \alpha_k(G_1)$, and if $S_1 = \emptyset$, then $|S| \leq \alpha_k(G_2)$. Suppose now that $S_1, S_2 \neq \emptyset$, and let $x = \Delta(G_1[S_1]) + 1$ and $y = \Delta(G_2[S_2]) + 1$. By the previous lemma, $x, y \in [k-1]$, $|S_1| \leq k - y$,

and $|S_2| \leq k - x$. Moreover, S_1 is x-independent and S_2 is y-independent, and hence $|S_1| \leq \alpha_x(G_1)$ and $|S_2| \leq \alpha_y(G_2)$. Therefore,

$$|S| \leq \min\{\alpha_x(G_1), k - y\} + \min\{\alpha_y(G_2), k - x\},$$

which proves the upper bound.

For the reverse inequality, clearly, $\alpha_k(G) \geq \alpha_k(G_1)$ and $\alpha_k(G) \geq \alpha_k(G_2)$. Now, fix $x, y \in [k-1]$ and let $a = \min\{\alpha_x(G_1), k - y\}$ and $b = \min\{\alpha_y(G_2), k - x\}$. Choose an x-independent set $T_1 \subseteq V(G_1)$ of cardinality a and a y-independent set $T_2 \subseteq V(G_2)$ of cardinality b. For every $v \in T_1$, we have $d_{T_1 \cup T_2}(v) \leq (x-1) + b \leq (x-1) + (k-x) = k-1$. Similarly, for every $v \in T_2$, we have $d_{T_1 \cup T_2}(v) \leq (y-1) + a \leq (y-1) + (k-y) = k-1$. Thus, $T_1 \cup T_2$ is k-independent, and consequently $\alpha_k(G) \geq a + b$. Since this holds for every $x, y \in [k-1]$, the result follows. □

The previous theorem also yields the following useful shortcut.

Corollary 8. *Let $k \geq 2$ and let $G = G_1 \otimes G_2$. If $\max\{\alpha_k(G_1), \alpha_k(G_2)\} \geq 2(k-1)$, then $\alpha_k(G) = \max\{\alpha_k(G_1), \alpha_k(G_2)\}$.*

Proof. By the previous corollary, every k-independent set containing vertices from both G_1 and G_2 has cardinality at most $2(k-1)$. Hence, under the hypothesis, no such set can have cardinality larger than $\max\{\alpha_k(G_1), \alpha_k(G_2)\}$. Since a maximum k-independent set of either G_1 or G_2 is also k-independent in G, the result follows. □

4. k-Independence in P_4-tidy graphs

In Section 3, we analyzed disjoint union and join; it remains to treat spiders and quasi-spiders. Since P_4-tidy graphs form a hereditary class and duplication affects only $C \cup X$, the head R always induces a P_4-tidy subgraph, so each structure can be handled recursively in terms of the values $\alpha_j(G[R])$. All results below follow the same strategy. First, a normalization lemma shows that some maximum k-independent set contains the independent side of the partition (Lemma 10; Lemma 14 for the types (R, C, X^t), where the twin edge requires $k \geq 2$; and Lemmas 16 and 18 for the types (R, C^f, X), the only ones with two canonical forms); all four rest on a single exchange principle, which we isolate in Lemma 9. Then degree-counting bounds, matched by explicit constructions, determine how many vertices of $R \cup C$ can complete the canonical set—each term of the resulting max formulas corresponds to one completion. Table 1 collects all formulas.

Lemma 9 (Exchange lemma)**.** *Let S be a k-independent set of a graph G, and let $c \in S$ and $x \in V(G) \setminus S$ be such that $N(x) \cap (S \setminus \{c\}) \subseteq N(c)$. Then $S' = (S \setminus \{c\}) \cup \{x\}$ is k-independent.*

Proof. Let $v \in S' \setminus \{x\}$. If $v \notin N(x)$, then v gains no neighbor and $d_{S'}(v) \leq d_S(v)$. If $v \in N(x)$, then $v \in N(c)$ by hypothesis, so v loses c and gains x, and again $d_{S'}(v) \leq d_S(v)$. Finally, $N(x) \cap S' \subseteq N(c) \cap (S \setminus \{c\})$, so $d_{S'}(x) \leq d_S(c) \leq k-1$. □

Lemma 10. *Let $k \geq 1$, and let G be either a spider or a quasi-spider of one of the types (R, C, X^f) or (R, C^t, X), thin or thick. Let $X^\star$ denote the independent side of G, that is, $X^\star = X$ for spiders and quasi-spiders of type (R, C^t, X), and $X^\star = X^f$ for quasi-spiders of type (R, C, X^f). Then there exists a maximum k-independent set S of G such that $X^\star \subseteq S$.*

Proof. Among all maximum k-independent sets of G, choose S maximizing $|S \cap X^\star|$, and suppose $x_i \in X^\star \setminus S$; in each case below we obtain a maximum k-independent set with more vertices of $X^\star$, contradicting the choice of S.

Suppose first that G is a thin spider or a thin quasi-spider of type (R, C, X^f), so that $N(x_i) = \{c_i\}$ (both false twins inherit c_s as their unique neighbor). If $c_i \notin S$, then x_i can be added to S; otherwise Lemma 9 applies to (c_i, x_i), the containment being vacuous.

Suppose next that G is a thin quasi-spider of type (R, C^t, X). If $i \neq s$, the previous argument applies. If $i = s$, then $N(x_s) = \{c_s, c'_s\}$: if neither twin is in S, then x_s can be added; otherwise pick $c \in S \cap \{c_s, c'_s\}$ and apply Lemma 9 to (c, x_s), the containment holding because the true twins are adjacent.

Finally, suppose that G is thick. Then every neighbor of x_i on the clique side is adjacent to all other clique-side vertices, so if $S \cap N(x_i) = \emptyset$, then x_i can be added, and otherwise Lemma 9 applies to (c_j, x_i) for any $c_j \in S \cap N(x_i)$, the containment holding because the clique side is a clique. Hence $X^\star \subseteq S$. □

Theorem 11. *Let $k \geq 1$, and let G be a spider with partition (R, C, X). Set $m = |C| = |X|$ and $r = |R|$.*

(i) If G is thin, then $\alpha_k(G) = m + \max\{\alpha_k(G[R]), \min\{k-1, m+r\}\}$.
(ii) If G is thick, then

$$\alpha_k(G) = \begin{cases} \min\{k+1, 2m+r\}, & \text{if } m \leq k \text{ and } r < k-m+1; \\ m + \alpha_k(G[R]), & \text{otherwise.} \end{cases}$$

Proof. By Lemma 10, in both cases we may consider a maximum k-independent set S such that $X \subseteq S$.

For (i), since R is anticomplete to X, if T is a maximum k-independent set of $G[R]$, then $X \cup T$ is k-independent, yielding $\alpha_k(G) \geq m + \alpha_k(G[R])$. Also, letting $q = \min\{k-1, m+r\}$, if $q \leq r$ we may choose any q vertices of R, while if $q > r$ we choose all vertices of R and $q - r$ vertices of C; together with X, these vertices form a k-independent set. Indeed, each selected vertex of C has one neighbor in X and $q-1$ among the other selected vertices, hence degree $q \leq k-1$; each selected vertex of R has degree at most $q-1$; and each vertex of X has degree at most one, which is allowed since a vertex of C is selected only if $q \geq 1$, that is, $k \geq 2$. Hence $\alpha_k(G) \geq m + \min\{k-1, m+r\}$. Conversely, if $S \cap C = \emptyset$, then $S \setminus X \subseteq R$ is k-independent in $G[R]$, so $|S| \leq m + \alpha_k(G[R])$. If $S \cap C \neq \emptyset$, take $c \in S \cap C$. Since c is adjacent to every vertex of $S \cap R$, every other vertex of $S \cap C$, and its matched vertex in X, we have

$$d_S(c) = |S \cap R| + |S \cap C| \leq k-1,$$

and consequently, $|S| \leq m + \min\{k-1, m+r\}$. This proves (i).

For (ii), the set $X \cup T$, where T is a maximum k-independent set of $G[R]$, again gives $\alpha_k(G) \geq m + \alpha_k(G[R])$. Suppose first that $m \leq k$ and $r < k-m+1$. Let $q = \min\{k-m+1, r+m\}$. Choosing all vertices of R and $q-r$ vertices of C, together with X, yields a k-independent set of cardinality $m+q = \min\{k+1, 2m+r\}$. Indeed, each selected vertex of C has $m-1$ neighbors in X and $q-1$ among the other selected vertices, hence degree $m+q-2 \leq k-1$; each vertex of R has degree at most $q-1$; and each vertex of X has at most $q-r \leq k-m+1 \leq k-1$ selected neighbors, since $m \geq 2$. For the upper bound, if $S \cap C = \emptyset$, then $|S| \leq m + \alpha_k(G[R])$. Otherwise, for any $c \in S \cap C$,

$$d_S(c) = |S \cap R| + |S \cap C| + m - 2 \leq k-1,$$

so $|S \cap R| + |S \cap C| \leq k-m+1$. Thus, $S \cap C \neq \emptyset$ is possible only when $m \leq k$, and in this case $|S| \leq \min\{k+1, 2m+r\}$. If moreover $r < k-m+1$, then $\alpha_k(G[R]) = r$ and $m + r \leq \min\{k+1, 2m+r\}$,

proving the first case. Otherwise, if $m > k$, necessarily $S \cap C = \emptyset$; and if $m \leq k$ and $r \geq k - m + 1$, then $\alpha_k(G[R]) \geq k - m + 1$, whence $k + 1 \leq m + \alpha_k(G[R])$. Therefore, every solution intersecting C is no larger than $m + \alpha_k(G[R])$, proving the second case. □

For quasi-spiders with a duplicated vertex in X, the analysis splits according to whether c_s enters the solution, since c_s is the unique clique vertex whose relation to the twin pair differs from that of the others. In the thick types a dichotomy recurs: when $k < m$, no clique-side vertex fits alongside the independent side, as its neighbors there already exhaust the degree budget.

Theorem 12. *Let $k \geq 1$, and let G be a quasi-spider of type (R, C, X^f), obtained by adding a false twin to $x_s \in X$. Set $m = |C| = |X|$ and $r = |R|$.*

(i) If G is thin, then $\alpha_1(G) = m + 1 + \alpha_1(G[R])$, and, for $k \geq 2$,

$$\alpha_k(G) = m + 1 + \max\{\alpha_k(G[R]),\ \min\{k - 1, r + m - 1\},\ \min\{k - 2, r + m\}\}.$$

(ii) If G is thick, then $\alpha_k(G) = m + 1 + \alpha_k(G[R])$ for $k < m$, while, for $k \geq m$,

$$\alpha_k(G) = m + 1 + \max\{\alpha_k(G[R]),\ \min\{k - m, r + m\},\ \min\{k - m + 1, r + 1\}\}.$$

Proof. By Lemma 10, in both cases we may consider a maximum k-independent set S such that $X^f \subseteq S$. Notice that $|X^f| = m + 1$.

For (i), if $k = 1$, the set X^f is independent and anticomplete to R, while every vertex of C has a neighbor in X^f. Hence $\alpha_1(G) = m + 1 + \alpha_1(G[R])$. Assume now that $k \geq 2$. If T is a maximum k-independent set of $G[R]$, then $X^f \cup T$ is k-independent.

For the second candidate, let $q = \min\{k-1, r+m-1\}$. If $q \leq r$, choose any q vertices of R; otherwise, choose all vertices of R and $q - r$ vertices of $C \setminus \{c_s\}$. Together with X^f, they form a k-independent set. Indeed, each selected c_j, $j \neq s$, has one neighbor in X^f and $q - 1$ among the other selected vertices, hence degree $q \leq k - 1$; each selected vertex of R has degree at most $q - 1$; and each vertex of X^f has degree at most $1 \leq k - 1$.

For the third candidate, let $q = \min\{k - 2, r + m\}$. If $q = 0$, take $Q = \emptyset$, in which case this candidate reduces to the baseline set X^f. Otherwise, choose $Q \subseteq R \cup C$ with $|Q| = q$ and $c_s \in Q$. Then $X^f \cup Q$ is k-independent: the vertex c_s has degree at most $q + 1 \leq k - 1$, while every other selected vertex has degree at most $q \leq k - 1$. Every vertex of X^f has degree at most one, since x_s and x'_s are adjacent only to c_s.

For the upper bound, if $S \cap C = \emptyset$, then $|S| \leq m + 1 + \alpha_k(G[R])$. If $S \cap C \neq \emptyset$ and $c_s \notin S$, any $c \in S \cap C$ has exactly one neighbor in X^f, and therefore $d_S(c) = |S \cap R| + |S \cap C| \leq k - 1$.

Since at most $r + m - 1$ vertices are available in $R \cup (C \setminus \{c_s\})$, we obtain

$$|S| \leq m + 1 + \min\{k - 1, r + m - 1\}.$$

Finally, if $c_s \in S$, then c_s is adjacent to both false twins and hence $d_S(c_s) = |S \cap R| + |S \cap C| + 1 \leq k - 1$, which yields $|S| \leq m + 1 + \min\{k - 2, r + m\}$. This proves (i).

For (ii), $X^f \cup T$ is again k-independent for every maximum k-independent set T of $G[R]$, giving $\alpha_k(G) \geq m + 1 + \alpha_k(G[R])$. If $k < m$, no vertex of C can belong to S: the vertex c_s has $m - 1$ neighbors in X^f, while every c_j, $j \neq s$, has m neighbors in X^f. Thus $S \cap C = \emptyset$ and $\alpha_k(G) = m + 1 + \alpha_k(G[R])$.

Suppose now that $k \geq m$. For the second candidate, let $q = \min\{k - m, r + m\}$. If $q = 0$, take $Q = \emptyset$. Otherwise, choose any $Q \subseteq R \cup C$ with $|Q| = q$. Then $X^f \cup Q$ is k-independent: every selected c_j, $j \neq s$, has m neighbors in X^f and at most $q - 1$ in Q, hence degree at most $m + q - 1 \leq k - 1$, while c_s, if selected, has one fewer neighbor in X^f; every vertex of $Q \cap R$ has degree at most $q - 1$; and every vertex of X^f has at most $|Q \cap C| \leq q \leq k - m \leq k - 1$ selected neighbors. This yields the second term.

For the third candidate, let $q = \min\{k - m + 1, r + 1\}$ and choose $Q \subseteq R$ with $|Q| = q - 1$. Then $X^f \cup Q \cup \{c_s\}$ is k-independent: c_s is nonadjacent to both false twins, so it has $m - 1$ neighbors in X^f and $q - 1$ in Q, hence degree $m + q - 2 \leq k - 1$; every vertex of Q has degree at most $q - 1$; and every vertex of X^f has degree at most one.

For the upper bound, if $S \cap C = \emptyset$, then $|S| \leq m + 1 + \alpha_k(G[R])$. If S contains some c_j with $j \neq s$, then c_j is adjacent to m vertices of X^f, every vertex of $S \cap R$, and every other vertex of $S \cap C$, so

$$m + |S \cap R| + |S \cap C| - 1 \leq k - 1.$$

Hence, $|S \cap R| + |S \cap C| \leq k - m$, and therefore $|S| \leq m + 1 + \min\{k - m, r + m\}$.

The only remaining possibility is $S \cap C = \{c_s\}$. Since c_s has exactly $m - 1$ neighbors in X^f, we obtain $|S \cap R| \leq k - m$, and therefore,

$$|S| \leq m + 1 + \min\{k - m + 1, r + 1\}.$$

Combining the three possibilities proves (ii). □

Theorem 13. *Let $k \geq 1$, and let G be a quasi-spider of type (R, C^t, X), obtained by adding a true twin to $c_s \in C$. Set $m = |C| = |X|$ and $r = |R|$.*

(i) If G is thin, then $\alpha_k(G) = m + \max\{\alpha_k(G[R]), \min\{k - 1, r + m + 1\}\}$.
(ii) If G is thick, then $\alpha_k(G) = m + \alpha_k(G[R])$ for $k < m$, while, for $k \geq m$,

$$\alpha_k(G) = m + \max\{\alpha_k(G[R]), \min\{k - m + 1, r + m + 1\}\}.$$

Proof. By Lemma 10, in both cases we may consider a maximum k-independent set S such that $X \subseteq S$.

For (i), since R is anticomplete to X, if T is a maximum k-independent set of $G[R]$, then $X \cup T$ is k-independent, giving $\alpha_k(G) \geq m + \alpha_k(G[R])$. Also, let $q = \min\{k - 1, r + m + 1\}$, and choose any $Q \subseteq R \cup C^t$ with $|Q| = q$. Then $X \cup Q$ is k-independent: every vertex of $Q \cap C^t$, including either true twin, has one neighbor in X and $q - 1$ in Q, hence degree $q \leq k - 1$; every vertex of $Q \cap R$ has degree at most $q - 1$; every x_i with $i \neq s$ has degree at most one; and x_s has at most $\min\{2, q\} \leq k - 1$ selected neighbors. Hence $\alpha_k(G) \geq m + \min\{k - 1, r + m + 1\}$. Conversely, if $S \cap C^t = \emptyset$, then $|S| \leq m + \alpha_k(G[R])$. If $S \cap C^t \neq \emptyset$, take $c \in S \cap C^t$. Since C^t is a clique and every vertex of C^t is complete to R, while each such vertex has at least one neighbor in X, we have $d_S(c) \geq |S \cap R| + |S \cap C^t|$, and therefore $|S \cap R| + |S \cap C^t| \leq k - 1$. Since $|R \cup C^t| = r + m + 1$, it follows that $|S| \leq m + \min\{k - 1, r + m + 1\}$, proving (i).

For (ii), $X \cup T$ again gives $\alpha_k(G) \geq m + \alpha_k(G[R])$. If $S \cap C^t \neq \emptyset$, every selected vertex of C^t has at least $m - 1$ neighbors in X, in addition to all selected vertices of R and all other selected vertices of C^t. Thus $d_S(c) \geq m - 1 + |S \cap R| + |S \cap C^t| - 1$, so $|S \cap R| + |S \cap C^t| \leq k - m + 1$. In particular, if $k < m$, then $S \cap C^t = \emptyset$, and consequently $\alpha_k(G) = m + \alpha_k(G[R])$. Suppose now that $k \geq m$, and let $q = \min\{k - m + 1, r + m + 1\}$. Choose any set $Q \subseteq R \cup C^t$ with $|Q| = q$, and write $a = |Q \cap C^t|$

and $b = |Q \cap R|$. For every $c \in Q \cap C^t$, $d_{X\cup Q}(c) = (m-1)+(a-1)+b = m+q-2 \le k-1$. Moreover, every $x \in X$ has at most $a \le q \le k-1$ neighbors in Q, while every $v \in Q \cap R$ satisfies $d_{X\cup Q}(v) \le a+b-1 = q-1 \le k-1$.

Hence, $X \cup Q$ is k-independent, and therefore $\alpha_k(G) \ge m + \min\{k-m+1, r+m+1\}$.

Conversely, if $S \cap C^t = \emptyset$, then $|S| \le m+\alpha_k(G[R])$. If $S \cap C^t \ne \emptyset$, the inequality obtained above gives $|S \cap R| + |S \cap C^t| \le k-m+1$. Since $|R \cup C^t| = r+m+1$, it follows that $|S| \le m + \min\{k-m+1, r+m+1\}$.

Combining the two possibilities proves (ii). □

When a vertex of X is duplicated into true twins, the edge inside the pair obstructs the normalization at $k = 1$; the next lemma shows this is the only obstruction.

Lemma 14. *Let $k \ge 2$, and let G be a thin or thick quasi-spider of type (R, C, X^t). Then there exists a maximum k-independent set S of G such that $X^t \subseteq S$.*

Proof. Choose a maximum k-independent set S maximizing $|S \cap X^t|$, and suppose $x_i \in X^t \setminus S$; in each case below we obtain a maximum k-independent set with more vertices of X^t, a contradiction.

If G is thin and $i \ne s$, then $N(x_i) = \{c_i\}$: if $c_i \notin S$, then $S \cup \{x_i\}$ is k-independent; otherwise, Lemma 9 applies to (c_i, x_i), the containment being vacuous. If G is thin and x_i is a twin, say x_s, then $N(x_s) = \{c_s, x'_s\}$: if $S \cap N(x_s) = \emptyset$, add x_s; if $c_s \notin S$ and $x'_s \in S$, then $d_S(x'_s) = 0$ and x_s can be added, since $k \ge 2$; if $c_s \in S$, Lemma 9 applies to (c_s, x_s), since $x'_s \in N(c_s)$.

If G is thick and $i \ne s$, then $N(x_i) = C \setminus \{c_i\}$: if $S \cap N(x_i) = \emptyset$, add x_i; otherwise, pick $c_j \in S \cap N(x_i)$ and apply Lemma 9 to (c_j, x_i), the containment holding because C is a clique. If G is thick and the missing vertex is x_s, then $N(x_s) = (C \setminus \{c_s\}) \cup \{x'_s\}$: if $S \cap (C \setminus \{c_s\}) = \emptyset$, then $d_S(x'_s) = 0$ and x_s can be added, since $k \ge 2$; otherwise, pick $c_j \in S \cap (C \setminus \{c_s\})$ and apply Lemma 9 to (c_j, x_s), since C is a clique and x'_s is adjacent to every vertex of $C \setminus \{c_s\}$. Hence, $X^t \subseteq S$. □

Theorem 15. *Let $k \ge 1$, and let G be a quasi-spider of type (R, C, X^t), obtained by adding a true twin to $x_s \in X$. Set $m = |C| = |X|$ and $r = |R|$.*

(i) If G is thin, then $\alpha_1(G) = m + \alpha_1(G[R])$, while, for $k \ge 2$,

$$\alpha_k(G) = m + 1 + \max\{\alpha_k(G[R]),\ \min\{k-1, r+m-1\},\ \min\{k-2, r+m\}\}.$$

(ii) If G is thick, then $\alpha_1(G) = m + \alpha_1(G[R])$ and $\alpha_k(G) = m + 1 + \alpha_k(G[R])$ for $2 \le k < m$, while, for $k \ge m$,

$$\alpha_k(G) = m + 1 + \max\{\alpha_k(G[R]),\ \min\{k-m, r+m\},\ \min\{k-m+1, r+1\}\}.$$

Proof. For $k = 1$, the two true twins cannot both belong to an independent set. By the usual exchange argument, every selected vertex of C can be replaced by a corresponding vertex of X^t, so there is a maximum independent set containing m vertices of X^t and no vertex of C. Since R is anticomplete to X^t, in both the thin and thick cases, $\alpha_1(G) = m + \alpha_1(G[R])$.

Assume henceforth that $k \ge 2$. By Lemma 14, we may consider a maximum k-independent set S such that $X^t \subseteq S$.

For (i), if T is a maximum k-independent set of $G[R]$, then $X^t \cup T$ is k-independent, yielding $\alpha_k(G) \ge m + 1 + \alpha_k(G[R])$. Let $q = \min\{k-1, r+m-1\}$. Choosing any set Q of q vertices of $R \cup (C \setminus \{c_s\})$ and adding it to X^t gives a k-independent set: each selected c_j has degree $q \le k-1$, each selected vertex of

R has degree at most $q-1$, each x_i with $i \neq s$ has degree at most one, and each true twin has degree one, its only neighbor in the set being the other twin; all these degrees are allowed since $k \geq 2$. For the third term, let $q = \min\{k-2, r+m\}$. If $q > 0$, choose $Q \subseteq R \cup C$ with $|Q| = q$ and $c_s \in Q$. Then c_s has degree $q+1 \leq k-1$, each other selected vertex of C has degree at most $q \leq k-1$, and the two true twins have degree two; the latter is allowed since $q > 0$ implies $k \geq 3$. Each selected vertex of R has degree at most $q-1$, and each x_i with $i \neq s$ has degree at most one. Thus $X^t \cup Q$ is k-independent.

Conversely, if $S \cap C = \emptyset$, then $|S| \leq m+1+\alpha_k(G[R])$. If $S \cap C \neq \emptyset$ and $c_s \notin S$, every $c \in S \cap C$ has one neighbor in X^t, so $d_S(c) = |S \cap R| + |S \cap C| \leq k-1$, and therefore $|S| \leq m+1+\min\{k-1, r+m-1\}$. If $c_s \in S$, then c_s is adjacent to both true twins, and hence $d_S(c_s) = |S \cap R| + |S \cap C| + 1 \leq k-1$, which gives $|S| \leq m+1+\min\{k-2, r+m\}$. This proves (i).

For (ii), $X^t \cup T$ again gives $\alpha_k(G) \geq m+1+\alpha_k(G[R])$. If $k < m$, no vertex of C can belong to S: c_s has $m-1$ neighbors in X^t, whereas every c_j, $j \neq s$, has m such neighbors. Hence, for $2 \leq k < m$, $\alpha_k(G) = m+1+\alpha_k(G[R])$.

Suppose now that $k \geq m$. Let $q = \min\{k-m, r+m\}$ and choose any $Q \subseteq R \cup C$ with $|Q| = q$. For every $c_j \in Q \cap (C \setminus \{c_s\})$, $d_{X^t \cup Q}(c_j) = m+q-1 \leq k-1$, while c_s, if selected, has one fewer neighbor in X^t. Every vertex of $Q \cap R$ has degree at most $q-1$, and every vertex of X^t has degree at most $q+1$. Since $q \leq k-m$ and $m \geq 2$, we have $q+1 \leq k-1$. Thus $X^t \cup Q$ is k-independent and gives the second term. For the third, let $q = \min\{k-m+1, r+1\}$ and choose $Q \subseteq R$ with $|Q| = q-1$. Then $X^t \cup Q \cup \{c_s\}$ is k-independent: c_s has $m-1$ neighbors in X^t and $q-1$ in Q, hence degree $m+q-2 \leq k-1$; every vertex of Q has degree at most $q-1$; every x_i with $i \neq s$ has degree one; and each true twin has degree one, since neither is adjacent to c_s.

For the upper bound, if $S \cap C = \emptyset$, then $|S| \leq m+1+\alpha_k(G[R])$. If S contains some $c_j \neq c_s$, then $m + |S \cap R| + |S \cap C| - 1 \leq k-1$, so $|S| \leq m+1+\min\{k-m, r+m\}$.

The only remaining possibility is $S \cap C = \{c_s\}$. Since c_s has $m-1$ neighbors in X^t, we obtain $|S \cap R| \leq k-m$, and consequently $|S| \leq m+1+\min\{k-m+1, r+1\}$.

Combining the three possibilities proves (ii). □

The types (R, C^f, X) are the only ones where the independent side may profitably be incomplete: the vertex x_s can be traded for the false-twin pair, which forces a second canonical form.

Lemma 16. *Let $k \geq 1$, and let G be a thin quasi-spider of type (R, C^f, X), obtained by adding a false twin c'_s to $c_s \in C$. Then there exists a maximum k-independent set S satisfying one of the following:*

(i) $X \subseteq S$;
(ii) $(X \setminus \{x_s\}) \cup \{c_s, c'_s\} \subseteq S$.

Proof. Choose a maximum k-independent set S maximizing $|S \cap X|$. If $x_i \notin S$ with $i \neq s$, then $N(x_i) = \{c_i\}$: either x_i can be added, or Lemma 9 applies to (c_i, x_i); both contradict the choice of S, so $X \setminus \{x_s\} \subseteq S$. If $x_s \in S$, condition (i) holds. Otherwise, $N(x_s) = \{c_s, c'_s\}$: if neither twin is in S, then x_s can be added, and if exactly one is, say c_s, then Lemma 9 applies to (c_s, x_s), the containment being vacuous since $c'_s \notin S$; both cases contradict the choice of S. Hence, $\{c_s, c'_s\} \subseteq S$, and condition (ii) holds. □

Theorem 17. *Let $k \geq 1$, and let G be a thin quasi-spider of type (R, C^f, X), obtained by adding a false*

twin c'_s to $c_s \in C$. Set $m = |C| = |X|$ and $r = |R|$. Then

$$\alpha_k(G) = \begin{cases} m + \max\{\alpha_k(G[R]),\ 1\}, & k \le 2; \\ m + \max\{\alpha_k(G[R]),\ \min\{k-1,\ r+m+1\},\ \min\{k,\ r+2\}\}, & k \ge 3. \end{cases}$$

Proof. Let S be a maximum k-independent set of G. By Lemma 16, we may assume that either $X \subseteq S$ or $(X \setminus \{x_s\}) \cup \{c_s, c'_s\} \subseteq S$.

We first establish the lower bounds. Let T be a maximum k-independent set of $G[R]$. Since X is anticomplete to R, $X \cup T$ is k-independent, and hence $\alpha_k(G) \ge m + \alpha_k(G[R])$. Moreover, $(X \setminus \{x_s\}) \cup \{c_s, c'_s\}$ is an independent set of cardinality $m+1$, yielding $\alpha_k(G) \ge m + \max\{\alpha_k(G[R]), 1\}$ for $k \le 2$.

Now, suppose that $k \ge 3$ and let $q = \min\{k-1, r+m+1\}$. Choose any $Q \subseteq R \cup C^f$ with $|Q| = q$, and write $a = |Q \cap C^f|$ and $b = |Q \cap R|$, so that $a + b = q$. We claim that $X \cup Q$ is k-independent. Each ordinary vertex of $Q \cap C^f$ has at most $1 + (a-1) + b = q$ neighbors in $X \cup Q$, while either false twin, if selected, has degree at most q. Each vertex of $Q \cap R$ has degree at most $a + b - 1 = q - 1$, every vertex of $X \setminus \{x_s\}$ has degree at most 1, and x_s has degree at most 2. Since $q \le k-1$ and $k \ge 3$, all these degrees are at most $k-1$. Therefore,

$$\alpha_k(G) \ge m + \min\{k-1, r+m+1\}.$$

For the third candidate, let $b = \min\{k-2, r\}$ and choose any $B \subseteq R$ with $|B| = b$. Then $X \cup \{c_s, c'_s\} \cup B$ is k-independent. Indeed, each false twin has degree $b+1 \le k-1$, every vertex of B has degree at most $2 + (b-1) = b + 1 \le k-1$, the vertex x_s has degree $2 \le k-1$, and every vertex of $X \setminus \{x_s\}$ has degree zero. Hence

$$\alpha_k(G) \ge m + 2 + \min\{k-2, r\} = m + \min\{k, r+2\}.$$

We now prove the upper bounds. First, suppose that $X \subseteq S$. If $S \cap C^f = \emptyset$, then $S \cap R$ is a k-independent set of $G[R]$, and thus $|S| \le m + \alpha_k(G[R])$. If $S \cap C^f = \{c_s, c'_s\}$, then x_s is adjacent to both false twins, so $d_S(x_s) = 2$ and this case can occur only if $k \ge 3$. Let $b = |S \cap R|$. Each false twin is adjacent to x_s and to every vertex of $S \cap R$, so $b+1 \le k-1$, and hence, $|S| = m+2+b \le m+\min\{k, r+2\}$.

Suppose now that $S \cap C^f \ne \emptyset$ and $S \cap C^f \ne \{c_s, c'_s\}$. This includes, in particular, the cases $S \cap C^f = \{c_s\}$ and $S \cap C^f = \{c'_s\}$. Put $a = |S \cap C^f|$ and $b = |S \cap R|$. If an ordinary vertex of C^f belongs to S, it is adjacent to its corresponding vertex of X, to every other selected vertex of C^f, and to every vertex of $S \cap R$, giving $a + b \le k-1$. If instead $S \cap C^f$ consists of exactly one false twin, that twin is adjacent to x_s and to every vertex of $S \cap R$, and again $a + b \le k-1$. Consequently,

$$|S| \le m + \min\{k-1, r+m+1\}.$$

Finally, suppose that the second canonical form holds. We may assume that $x_s \notin S$, since otherwise $X \subseteq S$ and the first canonical form applies. Thus, $x_s \notin S$. If no other vertex of C^f belongs to S, let $b = |S \cap R|$. Each false twin is adjacent to every vertex of $S \cap R$, so $b \le k-1$. Thus, for $k \ge 3$, $|S| = m + 1 + b \le m + 1 + \min\{k-1, r\} \le m + \min\{k, r+2\}$. If $k \le 2$, necessarily, $b = 0$, since every vertex of $S \cap R$ would otherwise have degree at least 2; hence, $|S| = m+1$.

If some vertex of $C^f \setminus \{c_s, c'_s\}$ also belongs to S, that vertex is adjacent to every selected vertex of R, to all other selected vertices of C^f, and to its corresponding vertex of $X \setminus \{x_s\}$. Hence, $|S \cap R| + |S \cap C^f| \le k-1$, and therefore

$$|S| \le (m-1) + \min\{k-1, r+m+1\} \le m + \min\{k-1, r+m+1\}.$$

It remains to combine the bounds. If $k \le 2$, the case $S \cap C^f = \{c_s, c'_s\}$ with $X \subseteq S$ does not occur, the other cases of the first canonical form give $|S| \le m + \max\{\alpha_k(G[R]), \min\{k-1, r+m+1\}\} \le m + \max\{\alpha_k(G[R]), 1\}$, and the second canonical form gives $|S| \le m+1$. Together with the lower bound $\alpha_k(G) \ge m + \max\{\alpha_k(G[R]), 1\}$, this proves the formula for $k \le 2$. If $k \ge 3$, every case yields $|S| \le m + \max\{\alpha_k(G[R]), \min\{k-1, r+m+1\}, \min\{k, r+2\}\}$, which matches the three lower bounds. □

Lemma 18. *Let $k \ge 1$, and let G be a thick quasi-spider of type (R, C^f, X), obtained by adding a false twin c'_s to $c_s \in C$. Then there exists a maximum k-independent set S satisfying one of the following:*

(i) $X \subseteq S$;
(ii) $\{x_s, c_s, c'_s\} \subseteq S$ and $S \cap (C^f \setminus \{c_s, c'_s\}) = \emptyset$.

Proof. Choose a maximum k-independent set S maximizing $|S \cap X|$; call the vertices of $C^f \setminus \{c_s, c'_s\}$ *ordinary*. In each case below, Lemma 9 or a direct addition yields a maximum k-independent set with more vertices of X, a contradiction.

First, $x_s \in S$, otherwise, since $N(x_s)$ consists exactly of the ordinary vertices, maximality gives an ordinary $c_j \in S$, and Lemma 9 applies to (c_j, x_s), as c_j is adjacent to every other vertex of C^f.

Now, suppose $x_i \notin S$ for some $i \ne s$. If an ordinary $c_j \in S$ with $j \ne i$ exists, Lemma 9 applies to (c_j, x_i). If $c_i \in S$, Lemma 9 applies to (c_i, x_i)—note that $c_i \notin N(x_i)$, which the lemma permits—since $N(x_i) \cap (S \setminus \{c_i\}) \subseteq \{c_s, c'_s\} \subseteq N(c_i)$. Hence S contains no ordinary vertex. If neither twin is in S, then x_i has no neighbor in S and can be added; if exactly one twin is in S, say c_s, Lemma 9 applies to (c_s, x_i), the containment being vacuous. Therefore, $\{c_s, c'_s\} \subseteq S$ and $S \cap (C^f \setminus \{c_s, c'_s\}) = \emptyset$; together with $x_s \in S$, condition (ii) holds. If no vertex of X is missing, condition (i) holds. □

In the thick case, the second canonical form is where $\alpha_{k-2}(G[R])$ enters: both false twins are adjacent to every selected vertex of R, consuming two units of its degree budget.

Theorem 19. *Let $k \ge 1$, and let G be a thick quasi-spider of type (R, C^f, X), obtained by adding a false twin c'_s to $c_s \in C$. Set $m = |C| = |X|$ and $r = |R|$. Then*

$$\alpha_k(G) = \begin{cases} \max\{m + \alpha_k(G[R]),\ 3\}, & k \le 2; \\ \max\{m + \alpha_k(G[R]),\ 3 + \min\{k-1,\ m-1+\alpha_{k-2}(G[R])\}\}, & 3 \le k < m; \\ \max\left\{\begin{array}{l} m + \alpha_k(G[R]), \\ m + \min\{k-m+1,\ r+m+1\}, \\ m + 2 + \min\{r,\ k-m\}, \\ 3 + \min\{k-1,\ m-1+\alpha_{k-2}(G[R])\} \end{array}\right\}, & k \ge \max\{3, m\}. \end{cases}$$

Proof. By Lemma 18, we may consider a maximum k-independent set S satisfying one of the two canonical forms given there.

Suppose first that $X \subseteq S$. If $S \cap C^f = \emptyset$, then $S \setminus X \subseteq R$ is k-independent in $G[R]$, and therefore $|S| \le m + \alpha_k(G[R])$.

This bound is attained by taking X together with a maximum k-independent set of $G[R]$.

Assume now that S contains some ordinary vertex $c_i \in C^f \setminus \{c_s, c'_s\}$. Let $a = |S \cap C^f|$ and $b = |S \cap R|$. Since c_i has $m-1$ neighbors in X and is adjacent to every vertex of $S \cap R$ and every other vertex of

$S \cap C^f$, $d_S(c_i) = (m-1) + (a-1) + b \leq k-1$, and hence $a + b \leq k - m + 1$. Thus this case is possible only if $k \geq m$ and $|S| \leq m + \min\{k - m + 1, r + m + 1\}$.

Conversely, let $q = \min\{k - m + 1, r + m + 1\}$ and choose any set $Q \subseteq R \cup C^f$ with $|Q| = q$. Writing $a = |Q \cap C^f|$ and $b = |Q \cap R|$, every ordinary vertex $c_i \in Q \cap C^f$ satisfies

$$d_{X \cup Q}(c_i) \leq (m-1) + (a-1) + b = m + q - 2 \leq k - 1.$$

The same upper bound holds for c_s and c'_s, whenever selected, since the false twins are nonadjacent. Moreover, every vertex of $Q \cap R$ has degree at most $q - 1$, and every vertex of X has at most $a \leq q \leq k - m + 1 \leq k - 1$ selected neighbors. Hence, $X \cup Q$ is k-independent and attains this bound.

It remains to consider, within the first canonical form, $S \cap C^f \subseteq \{c_s, c'_s\}$. If exactly one twin is selected, that twin has $m - 1$ neighbors in X and is adjacent to every vertex of $S \cap R$, which gives $|S \cap R| \leq k - m$, and hence $|S| \leq m + 1 + \min\{r, k - m\}$; for $k \geq 3$ this is dominated by the bound below. If both are selected, each has $m - 1$ neighbors in X and is adjacent to all selected vertices of R. Thus $|S \cap R| \leq k - m$. Moreover, every x_i with $i \neq s$ is adjacent to both false twins and therefore has degree two in $G[S]$; since $m \geq 2$, such a vertex exists, and hence this case can occur only if $k \geq 3$. Consequently, for $k \geq m$ and $k \geq 3$, $|S| \leq m + 2 + \min\{r, k - m\}$.

This bound is attained by taking $X \cup \{c_s, c'_s\}$ together with any set $B \subseteq R$ with $|B| = \min\{r, k - m\}$. Indeed, each vertex of B has degree at most $2 + (|B| - 1) \leq k - m + 1 \leq k - 1$ in the resulting set; each false twin has $m - 1$ neighbors in X and $|B|$ in B, hence degree at most $(m - 1) + (k - m) = k - 1$; the vertex x_s has no selected neighbor; and every other vertex of X has degree $2 \leq k - 1$.

Consider now the second canonical form. Let $Y = S \cap (X \setminus \{x_s\})$ and $T = S \cap R$. Since $c_s, c'_s \in S$, each false twin is adjacent to every vertex of $Y \cup T$, and therefore, $|Y| + |T| \leq k - 1$.

Moreover, every vertex of T already has c_s and c'_s as neighbors. Thus $\Delta(G[T]) \leq k - 3$, so T is $(k-2)$-independent in $G[R]$ and $|T| \leq \alpha_{k-2}(G[R])$. Since $|Y| \leq m - 1$, for $k \geq 3$, we obtain

$$|S| \leq 3 + \min\{k - 1, m - 1 + \alpha_{k-2}(G[R])\}.$$

To see that this bound is attained, let $q = \min\{k - 1, m - 1 + \alpha_{k-2}(G[R])\}$, and let T^* be a maximum $(k-2)$-independent set of $G[R]$. Choose $T \subseteq T^*$ with $|T| = \min\{q, \alpha_{k-2}(G[R])\}$, and then choose $Y \subseteq X \setminus \{x_s\}$ with $|Y| = q - |T|$. Such a set Y exists since $q \leq m - 1 + \alpha_{k-2}(G[R])$. Then $\{x_s, c_s, c'_s\} \cup Y \cup T$ is k-independent: each false twin has degree $q \leq k - 1$, the vertex x_s has no selected neighbor, every vertex of Y has degree two, and every vertex of T has the two false twins as neighbors and at most $k - 3$ neighbors in T. Hence, the fourth bound is attained.

Suppose that $k \leq 2$. In the second canonical form, every vertex of $Y \cup T$ would be adjacent to both false twins and have degree at least two, so $Y = T = \emptyset$ and $|S| = 3$; this value is attained, since $\{x_s, c_s, c'_s\}$ is an independent set. In the first canonical form with $S \cap C^f \neq \emptyset$, three cases arise. If S contains an ordinary vertex, then $a + b \leq k - m + 1$ with $a \geq 1$ forces $m = k = 2$ and $a + b = 1$, so $|S| = m + 1 = 3$. If $S \cap C^f$ consists of exactly one false twin, then $|S \cap R| \leq k - m$ again forces $m = k = 2$ and $S \cap R = \emptyset$, so $|S| = 3$. Finally, as shown above, $S \cap C^f = \{c_s, c'_s\}$ is impossible for $k \leq 2$. Therefore, $\alpha_k(G) = \max\{m + \alpha_k(G[R]), 3\}$.

If $3 \leq k < m$, no vertex of C^f can be selected together with all of X, since every vertex of C^f has $m - 1 \geq k$ neighbors in X. Thus only the first and fourth bounds remain, giving

$$\alpha_k(G) = \max\left\{m + \alpha_k(G[R]), 3 + \min\{k - 1, m - 1 + \alpha_{k-2}(G[R])\}\right\}.$$

Finally, if $k \geq \max\{3, m\}$, all four bounds are attainable, and taking their maximum gives the stated formula. □

Table 1 summarizes the results of this section, writing $a_j = \alpha_j(G[R])$.

Table 1. The k-independence number of spiders and quasi-spiders.

Structure	Type	$\alpha_k(G)$
Spider	Thin	$m + \max\{a_k, \min\{k-1, m+r\}\}$
	Thick	$\begin{cases} \min\{k+1, 2m+r\}, & m \leq k,\ r < k-m+1, \\ m + a_k, & \text{otherwise} \end{cases}$
Thin quasi	(R, C, X^f)	$\begin{cases} m+1+a_1, & k=1, \\ m+1+\max\{a_k, \min\{k-1, r+m-1\}, \min\{k-2, r+m\}\}, & k \geq 2 \end{cases}$
	(R, C, X^t)	$\begin{cases} m+a_1, & k=1, \\ m+1+\max\{a_k, \min\{k-1, r+m-1\}, \min\{k-2, r+m\}\}, & k \geq 2 \end{cases}$
	(R, C^t, X)	$m + \max\{a_k, \min\{k-1, r+m+1\}\}$
	(R, C^f, X)	$\begin{cases} m + \max\{a_k, 1\}, & k \leq 2, \\ m + \max\{a_k, \min\{k-1, r+m+1\}, \min\{k, r+2\}\}, & k \geq 3 \end{cases}$
Thick quasi	(R, C, X^f)	$\begin{cases} m+1+a_k, & k < m, \\ m+1+\max\{a_k, \min\{k-m, r+m\}, \min\{k-m+1, r+1\}\}, & k \geq m \end{cases}$
	(R, C, X^t)	$\begin{cases} m+a_1, & k=1, \\ m+1+a_k, & 2 \leq k < m, \\ m+1+\max\{a_k, \min\{k-m, r+m\}, \min\{k-m+1, r+1\}\}, & k \geq m \end{cases}$
	(R, C^t, X)	$\begin{cases} m+a_k, & k < m, \\ m+\max\{a_k, \min\{k-m+1, r+m+1\}\}, & k \geq m \end{cases}$
	(R, C^f, X)	$\begin{cases} \max\{m+a_k, 3\}, & k \leq 2, \\ \max\{m+a_k, 3+\min\{k-1, m-1+a_{k-2}\}\}, & 3 \leq k < m, \\ \max\begin{cases} m+a_k, & m+\min\{k-m+1, r+m+1\}, \\ m+2+\min\{r, k-m\}, \\ 3+\min\{k-1, m-1+a_{k-2}\} \end{cases}, & k \geq \max\{3, m\} \end{cases}$

We record some boundary-case checks for the formulas in Table 1. If the head is empty, all expressions remain valid under the convention $\alpha_j(\emptyset) = 0$. When $m = 2$, thin and thick spiders are isomorphic, and the corresponding formulas coincide: if $r < k-1$, the thin-spider expression reduces to $\min\{k+1, r+4\}$, while otherwise both expressions reduce to $2 + \alpha_k(G[R])$. In the true-twin cases, $k = 1$ is treated separately because the two true twins cannot both belong to an independent set, whereas the formulas for $k = 2$ allow both twins when their induced degree remains at most one; in particular, when $m = 2$, the thick case (R, C, X^t) with $k = 2$ falls into the branch $k \geq m$. Finally, the thresholds at $k = m$ are consistent with the appearance of new feasible configurations exactly at this value; in particular, for the thick quasi-spider (R, C^f, X), the adjacent branches give the same value at $k = m$

whenever $m \geq 3$.

As an additional consistency check, we computationally verified the complete procedure on small instances: the construction of the labelled decomposition tree of Lemma 21, the formulas of Table 1, and the dynamic program of Section 5. Test graphs are generated recursively from the structure theorem, so that the corpus contains thin and thick spiders of weights $m \in \{2, 3, 4\}$, all eight types of quasi-spiders, the three exceptional graphs, and heads of three kinds: empty, every simple graph on at most four vertices up to isomorphism, and structured heads that are themselves spiders, quasi-spiders, unions, and joins, so that decompositions are nested several levels deep. The generator is not relied upon: every graph produced is independently checked against the definition of a P_4-tidy graph, that is, every vertex set inducing a P_4 is verified to have at most one partner. For each graph, the labelled tree is built by the procedure of Lemma 21, the labels are compared with the parameters the graph was constructed with, and the value returned by Algorithm 2 is compared, for every $1 \leq k \leq |V(G)|$, with the exact value obtained by exhaustive enumeration of all vertex subsets. Over a corpus of 7,421 graphs of order at most 13, comprising 42,883 tree nodes, 83,904 comparisons were performed and no discrepancy was found. This computational verification is not used in the proofs but provides an independent check of the case distinctions, of the boundary values, and of the construction of the decomposition tree. The verification code, together with a description of the testing procedure, is publicly available at Zenodo [19].

5. The algorithm

We now combine the results of Sections 3 and 4 into a dynamic programming algorithm over the decomposition tree $T(G)$ of a P_4-tidy graph G. Recall that the nodes of $T(G)$ are single vertices, exceptional nodes (P_5, $\overline{P_5}$ and C_5), and union, join, spider, and quasi-spider nodes. For the exceptional nodes, the required values are obtained by direct inspection.

Observation 20. *For the exceptional graphs, we have*

$$\alpha_k(P_5) = \begin{cases} 3, & k = 1, \\ 4, & k = 2, \\ 5, & k \geq 3, \end{cases} \qquad \alpha_k(\overline{P_5}) = \begin{cases} 2, & k = 1, \\ 3, & k = 2, \\ 4, & k = 3, \\ 5, & k \geq 4, \end{cases} \qquad \alpha_k(C_5) = \begin{cases} 2, & k = 1, \\ 3, & k = 2, \\ 5, & k \geq 3. \end{cases}$$

Algorithm 1 implements the characterization of k-independent sets in a join given by Theorem 7. A solution may be entirely contained in one of the two operands, in which case its size is bounded by $\alpha_k(H_1)$ or $\alpha_k(H_2)$; this also covers the case $k = 1$, for which the loops are empty and line 1 returns $\max\{\alpha_1(H_1), \alpha_1(H_2)\}$, in accordance with the identity $\alpha_1(H_1 \otimes H_2) = \max\{\alpha_1(H_1), \alpha_1(H_2)\}$. Otherwise, the solution intersects both operands: for each pair of admissible internal degree bounds x and y, the algorithm determines how many vertices can be selected from each side while accounting for the additional neighbors created by the join, and takes the maximum over all such choices. Since there are $O(k^2)$ pairs (x, y), the recurrence is evaluated in $O(k^2)$ time once the values $\alpha_i(H_1)$ and $\alpha_i(H_2)$, $i \in [k]$, are available.

Algorithm 1: JOIN(H_1, H_2, k)

```
M ← max{α_k(H_1), α_k(H_2)};
for x ← 1 to k − 1 do
    for y ← 1 to k − 1 do
        M ← max{M, min{α_x(H_1), k − y} + min{α_y(H_2), k − x}};
return M;
```

Obtaining the labelled decomposition tree. A module of G is a set $M \subseteq V(G)$ such that every vertex outside M is adjacent either to all or to none of the vertices of M; a module is strong if it does not overlap any other module. Let $MD(G)$ denote the modular decomposition tree of G, which can be computed in $O(n + |E(G)|)$ time [20, 21]. Its leaves are the vertices of G, and each internal node N represents the strong module $M(N)$ formed by the leaves below it; the children of N are the maximal strong modules properly contained in $M(N)$. The node N is parallel if $G[M(N)]$ is disconnected, series if $\overline{G[M(N)]}$ is disconnected, and prime otherwise. For a prime node N, the representative graph Q_N is the subgraph of G induced by one vertex chosen from each child of N; it does not depend on the choice of the vertices, and it is a prime graph, that is, a graph whose only modules are trivial.

The structure theorem of [8] can be stated in terms of $MD(G)$ as follows: if G is P_4-tidy and N is a prime node, then either $Q_N \in \{C_5, P_5, \overline{P_5}\}$ and every child of N is a leaf, or Q_N is a spider whose head has at most one vertex. In the latter case, the child in the head position, if any, is the module R, and every other child is a leaf, except for at most one child that is a module with exactly two vertices, namely the twin pair of a quasi-spider. We refer to this statement as the structure theorem in the proof below.

Lemma 21. *Let G be a P_4-tidy graph with n vertices. A decomposition tree $T(G)$ with $O(n)$ nodes, in which every union node is labelled as such, every join node is binary, every exceptional node is labelled with its isomorphism type, and every spider or quasi-spider node is labelled with its thin/thick type, the integers m and r, a pointer to the child representing $G[R]$, and, when applicable, the side (C or X) and the kind (true or false) of the twin pair, can be computed in $O(n + |E(G)|)$ time.*

Proof. Compute $MD(G)$ and, by one traversal, the parent of each node, the number $n(N) = |M(N)|$ of leaves below each node N, and preorder intervals, so that the test $y \in M(N)$ takes $O(1)$ time.
Union and join nodes. Parallel nodes become union nodes. A series node with children $H_1, \ldots, H_\ell$ is replaced by a binary tree with $\ell - 1$ join nodes. Since the sum of the numbers of children over all nodes of $MD(G)$ is $O(n)$, this transformation introduces only $O(n)$ additional nodes.
Representative graphs. Let N be a prime node and let Q_N be its representative graph, whose vertices are the children of N. Call a child of N *small* if it has at most two vertices. By the structure theorem, at most one child of N is not small, and such a child, denoted H_N, occupies the head position. For each vertex y of G, let $s(y)$ be the small child of a prime node that contains y, if one exists: $s(y) = y$ if the parent of y is prime, and $s(y)$ is the parent of y if that parent has two vertices and its own parent is prime. These cases are mutually exclusive, so $s(y)$ is well defined, and it is computed for all vertices in $O(n)$ time.

To obtain the degrees in Q_N, choose a representative vertex v in each small child of N and scan its adjacency list. Because Q_N does not depend on the choice of representatives, the neighbors of a single

representative determine the neighbors of its child in Q_N. A neighbor y of v lies in the small child $s(y)$ of N if $s(y)$ is defined and is a child of N; it lies in H_N if this is not the case and $y \in M(N)$; and it lies outside $M(N)$ otherwise. Since the vertices of Q_N are the children of N, an edge of G joining two vertices of the same child must not contribute to the degree of that child. As the child containing v has at most two vertices, the only neighbor of v that can lie in it is its other vertex v', and this happens precisely when that child is a series node, in which case v and v' are true twins; if the child is a parallel node, then v and v' are nonadjacent and no such edge occurs. We therefore mark the child $s(v)$ before the scan begins, so that it is never counted among the distinct children encountered. Counting the remaining distinct children with a marking array gives the degree of the child containing v, and the degree in Q_N of the vertex corresponding to H_N, if it exists, is the number of small children adjacent to H_N, a count accumulated during the same scans. Each scan marks at most $d_G(v) + 1$ entries of the array, one per neighbor of v together with $s(v)$, and only those entries are reset before the next scan; hence the scan takes $O(d_G(v) + 1)$ time, and excluding the intra-module edge adds a constant to it. The representative of a small child is a vertex v whose value $s(v)$ is that child; since s is a function, each vertex is the representative of at most one small child, and the adjacency list of each vertex is therefore scanned at most once over all prime nodes. The total time spent on these scans is thus bounded by $O(\sum_{v \in V(G)}(d_G(v) + 1)) = O(n + |E(G)|)$.

Exceptional nodes. If N has five children, all of them leaves, the structure theorem restricts Q_N to C_5, P_5, $\overline{P_5}$, or the five-vertex spider (the bull). These four graphs are distinguished by their degree sequences, namely $(2, 2, 2, 2, 2)$, $(1, 1, 2, 2, 2)$, $(2, 2, 2, 3, 3)$, and $(1, 1, 2, 3, 3)$, respectively. In the first three cases, N is labelled as an exceptional node of the corresponding type.

Spider and quasi-spider nodes. Otherwise Q_N is a prime spider with p vertices, of weight $t = \lfloor p/2 \rfloor$ and with $h = p - 2t \in \{0, 1\}$ head vertices. If Q_N is thin, the vertices of X have degree one, their neighbors form C, and the possible remaining vertex is the head. If Q_N is thick and $t \geq 3$, the vertices of X, the possible head, and the vertices of C have degrees $t - 1$, t, and $2t - 2 + h$, respectively, and $t - 1 < t < 2t - 2 + h$. Hence, Q_N is thin if and only if it has a vertex of degree one, and in both cases the partition is read from the degrees. For $t = 2$, thin and thick spiders are isomorphic, and we use the thin label; the corresponding formulas in Table 1 coincide.

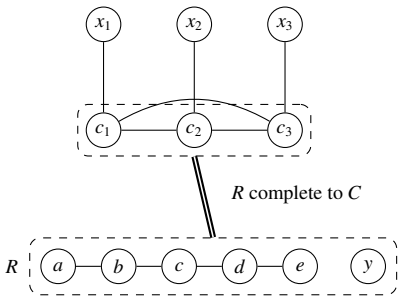
We store $m = t$ and $r = n(H)$, where H_N is the child occupying the head position, with $r = 0$ when the head is absent. If a two-vertex child occupies a position of C or X, then N is a quasi-spider node, and we record that side and whether the child is a series or a parallel node, corresponding respectively to true or false twins; otherwise N is a spider node. No index of the duplicated vertex needs to be stored, since the formulas of Table 1 depend only on the side containing the twin pair, its true/false kind, the thin/thick type, m, r, and the required values for $G[R]$.

The children corresponding to $C \cup X$ need not be retained as recursive children in the dynamic-programming representation, since their contribution is completely determined by the stored structural label. Only the child representing $G[R]$, when present, is recursively processed.

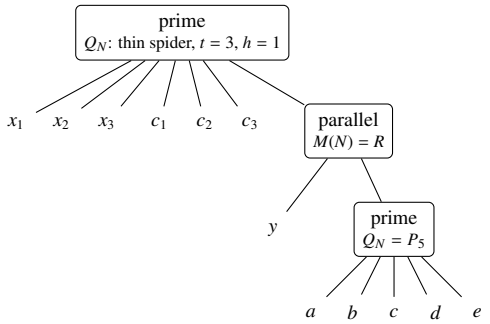
Thus all labels required by the dynamic program are obtained in $O(n+|E(G)|)$ time, and the resulting tree has $O(n)$ nodes. □

Let G be the thin spider of weight $m = 3$ whose head is $R = P_5 \oplus K_1$ (Figure 2(a)); G is P_4-tidy by Theorem 2. In $MD(G)$ (Figure 2(b)), the root is a prime node whose children are the six leaves $x_1, x_2, x_3, c_1, c_2, c_3$ and the parallel node representing R. Its representative graph Q_N is a thin spider of weight 3 with one head vertex: x_1, x_2, x_3 have degree one, their neighbors c_1, c_2, c_3 have degree

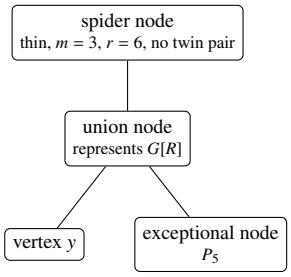
$t + h = 4$, and the head vertex has degree $t = 3$. Following the proof of Lemma 21, the root is labelled as a thin spider node with $m = 3$, $r = |R| = 6$ and no twin pair; the parallel node becomes a union node; and the prime node below it, whose five children are leaves with degree sequence $(1, 1, 2, 2, 2)$, becomes an exceptional node P_5. The resulting tree $T(G)$ is shown in Figure 2(c). The leaves x_i and c_i are not retained, since their contribution is determined by the label (m, r). Running Algorithm 2 on $T(G)$, Observation 20 gives $\alpha_i(P_5) = 3, 4, 5, 5, \ldots$, Observation 3 gives $\alpha_i(G[R]) = 4, 5, 6, 6, \ldots$, and the thin-spider formula of Table 1, $\alpha_i(G) = 3 + \max\{\alpha_i(G[R]), \min\{i-1, 9\}\}$, gives $\alpha_i(G) = 7, 8, 9, 9, \ldots$ for $i = 1, 2, 3, 4, \ldots$, which agrees with exhaustive enumeration.



(a) The thin spider G with head $R = P_5 \oplus K_1$.



(b) The modular decomposition tree $MD(G)$.



(c) The labelled decomposition tree $T(G)$.

Figure 2. A P_4-tidy graph G with nonempty head, its modular decomposition tree $MD(G)$, and the labelled decomposition tree $T(G)$ produced by Lemma 21. In (a), the double line indicates that R is complete to C.

By Lemma 21, the labelled tree $T(G)$ required by Algorithm 2 is obtained in $O(n + |E(G)|)$ time. Observe that the formulas of Table 1 depend only on the type of the node, on m and r, and on the values $\alpha_j(G[R])$.

For each node u of the decomposition tree, let G_u denote the subgraph represented by the subtree rooted at u. The dynamic programming table associated with u stores the values $\alpha_j(G_u)$ for $1 \le j \le k$. Thus, each node requires $O(k)$ entries and, since the decomposition tree has $O(n)$ nodes, the total state space is $O(kn)$.

For a spider or quasi-spider node, no optimization procedure is needed. The only formula in Table 1 that requires a value other than $\alpha_i(G[R])$ is that of the thick quasi-spider of type (R, C^f, X), which also uses $\alpha_{i-2}(G[R])$. Since the values are computed for $i = 1, 2, \ldots, k$ in increasing order, $\alpha_{i-2}(G[R])$

has already been computed and stored when it is needed. Hence every spider or quasi-spider node is processed in $O(1)$ time by table lookup. This is also the reason for the ordering of the loops in Algorithm 2: processing all nodes of $T(G)$ for each value of i in turn guarantees that, when a node is processed at level i, the values α_x of all its descendants, for every $x \le i$, have already been computed, as required by Algorithm 1.

Algorithm 2: k-independence number of a P_4-tidy graph

Input: A P_4-tidy graph G and an integer $k \ge 1$.
Output: $\alpha_k(G)$.

1 **if** $k \ge n(G)$ **then return** $n(G)$;
2 Construct the labelled decomposition tree $T(G)$;
3 **for** $i \leftarrow 1$ **to** k **do**
4 **foreach** *node H of $T(G)$ in bottom-up order* **do**
5 **if** $n(H) \le i$ **then**
6 $\alpha_i(H) \leftarrow n(H)$;
7 **else if** $H = H_1 \oplus \cdots \oplus H_\ell$ **then**
8 $\alpha_i(H) \leftarrow \sum_{j=1}^{\ell} \alpha_i(H_j)$;
9 **else if** $H = H_1 \otimes H_2$ **then**
10 $\alpha_i(H) \leftarrow \textsc{Join}(H_1, H_2, i)$;
11 **else if** *H is a spider or quasi-spider* **then**
12 compute $\alpha_i(H)$ using Table 1;
13 **else if** $H \in \{P_5, \overline{P_5}, C_5\}$ **then**
14 compute $\alpha_i(H)$ by Observation 20;
15 **return** $\alpha_k(G)$;

Theorem 22. *Let G be a P_4-tidy graph with n vertices, and let $k \ge 1$ be an integer. Algorithm 2 computes $\alpha_k(G)$ in $O(k^3 n + |E(G)|)$ time and uses $O(kn)$ space.*

Proof. Correctness. If $k \ge n(G)$, the value returned in line 1 is correct by Observation 4. Otherwise, we show, by induction on the bottom-up processing of $T(G)$, that after iteration i of the outer loop the value $\alpha_i(H)$ is correct for every node H of $T(G)$. If $n(H) \le i$, correctness follows from Observation 4. For union nodes, it follows from Observation 3. For a join node $H = H_1 \otimes H_2$, correctness follows from Theorem 7. Since the tree is processed bottom-up, the values $\alpha_i(H_1)$ and $\alpha_i(H_2)$ have already been computed in the current iteration, while every value $\alpha_x(H_j)$ with $x < i$ required by Algorithm 1 was computed in an earlier iteration of the outer loop.

For spider and quasi-spider nodes, correctness follows from Theorems 11–13, 15, 17, and 19, as summarized in Table 1. In particular, the value $\alpha_{i-2}(G[R])$ required for the thick quasi-spider of type (R, C^f, X) is available from iteration $i-2$; when $i \le 2$, the corresponding branch of the formula does not apply. Finally, for the exceptional nodes, correctness follows from Observation 20.

Complexity. Line 1 is a single comparison; if the algorithm continues, then $k < n$. The labelled decomposition tree $T(G)$ can be constructed in $O(n + |E(G)|)$ time, by Lemma 21, and has $O(n)$ nodes. For a fixed i, the total cost of all union nodes is $O(n)$, since the number of children summed over all nodes of $T(G)$ is $O(n)$. A join node costs $O(i^2)$ by Algorithm 1, while each spider, quasi-spider, or

exceptional node costs $O(1)$. Hence, over all $i = 1, \ldots, k$, the total cost of each join node is $\sum_{i=1}^{k} O(i^2) = O(k^3)$, and the dynamic programming phase takes $O(k^3 n)$ time. Therefore, the overall running time is $O(k^3 n + |E(G)|)$. Finally, storing $\alpha_i(H)$ for $1 \le i \le k$ at each node requires $O(kn)$ space. □

We remark that the running time is polynomial in n and k jointly, with k part of the input. For comparison, the metatheorem-based approaches discussed above treat k as a constant of the logical formula, while the generic dynamic programming algorithm over clique-width expressions [17] handles k as input but runs in $n^{O(w)}$ time; since $w \le 4$ on this class, that is a polynomial of fixed but high degree, whereas the algorithm above runs in linear time for every fixed k. Moreover, since the dynamic programming phase is executed only when $k < n$, by line 1 of Algorithm 2, the running time is $O(n^4)$ in the worst case.

6. Concluding remarks

We derived closed formulas for the k-independence number under disjoint union and join and for every type of spider and quasi-spider arising in the structural decomposition of P_4-tidy graphs, and combined them into a dynamic programming algorithm running in $O(k^3 n + |E(G)|)$ time. This extends the corresponding result for P_4-sparse graphs [13] to a substantially broader class of graphs with few induced P_4's.

Two directions seem natural for future work. First, in ongoing work we study the related notion of k-well-coveredness in this setting: a graph is k-well-covered if all its maximal k-independent sets have the same cardinality. For $k = 1$, well-coveredness of P_4-tidy graphs was completely characterized by Klein, de Mello, and Morgana [11]; extending that characterization to general k requires analyzing every maximal k-independent set, rather than one of maximum cardinality as in the normalization lemmas of Section 4. Second, the techniques developed here may apply to other graph classes admitting a decomposition into unions, joins, and bounded families of prime structures, such as P_4-laden graphs. More ambitiously, they may also apply to classes of bounded clique-width for which explicit formulas, rather than metatheorem-based algorithms, are of interest.

A further direction is the enumeration of k-independent sets by cardinality. Let $N_j(G, i)$ denote the number of j-independent sets of G of cardinality i, so that $\alpha_k(G) = \max\{i : N_k(G, i) > 0\}$. The disjoint union satisfies the convolution identity $N_k(G_1 \oplus G_2, i) = \sum_{i_1+i_2=i} N_k(G_1, i_1)\, N_k(G_2, i_2)$, that is, the corresponding generating polynomials $\sum_i N_k(G, i)\, x^i$ multiply. The join is more delicate: a set $S = S_1 \cup S_2$ with $S_\ell \subseteq V(G_\ell)$ is k-independent in $G_1 \otimes G_2$ if and only if every vertex of S_1 has at most $k - 1 - |S_2|$ neighbors in S_1 and every vertex of S_2 has at most $k - 1 - |S_1|$ neighbors in S_2. Thus the two sides are coupled through their cardinalities, and such a recurrence must carry the whole table $N_j(\cdot, i)$ rather than a single value. Whether the structures of the P_4-tidy decomposition admit closed counting analogues of the formulas of Table 1 seems to us the interesting question. The situation parallels the one discussed in Section 1: for every fixed k, tractability on graphs of bounded clique-width already follows from the counting version of the metatheorem [22], so the contribution would again lie in explicit formulas and in an explicit dependence on k.

Author contributions

Márcia R. Cappelle: Conceptualization, Methodology, Validation, Writing–review and editing, Supervision; Luan César D. Carvalho: Methodology, Software, Investigation, Writing–original draft, Writing–review and editing. All authors have read and approved the final version of the manuscript for publication.

Use of Generative-AI tools declaration

The authors used generative AI tools—Claude (Anthropic) and ChatGPT (OpenAI)—as assistive tools in the preparation of this work: for language editing and for drafting and revising portions of the text; for producing the figures; for writing the code of the computational verification deposited at [19]; and for cross-checking proofs, case distinctions, and boundary values against independently computed values. All definitions, statements, and proofs are the authors' own. The authors have reviewed and validated every AI-assisted output used in this work, including the computational verification code, and take full responsibility for the content of this article.

Acknowledgments

The authors thank the anonymous reviewers for their careful reading and constructive comments, which helped improve the presentation and technical accuracy of the manuscript. This research was supported by the Fundação de Amparo à Pesquisa do Estado de Goiás (FAPEG), under Public Call No. 21/2024, Process No. 202510267000505.

Conflict of interest

All authors declare no conflicts of interest in this paper.

References

1. J. F. Fink, M. S. Jacobson, *n*-Domination in graphs, In: *Graph theory with applications to algorithms and computer science*, New York: John Wiley and Sons, 1985, 282–300.
2. M. S. Jacobson, K. Peters, Complexity questions for *n*-domination and related parameters, *Proceedings of the 18th Manitoba Conference on Numerical Mathematics and Computing*, Winnipeg, Canada, 1989, 7–22.
3. M. Blidia, M. Chellali, O. Favaron, N. Meddah, Maximal *k*-independent sets in graphs, *Discuss. Math. Graph T.*, **28** (2008), 151–163. https://doi.org/10.7151/dmgt.1398
4. Y. Caro, A. Hansberg, New approach to the *k*-independence number of a graph, *Electron. J. Comb.*, **20** (2013), #P33. https://doi.org/10.37236/2646
5. M. Chellali, O. Favaron, A. Hansberg, L. Volkmann, *k*-Domination and *k*-independence in graphs: A survey, *Graph. Combinator.*, **28** (2012), 1–55. https://doi.org/10.1007/s00373-011-1040-3

6. M. Cappelle, E. Coelho, O. Mortosa, J. Nascimento, Complexity results on k-independence in some graph products, *RAIRO-Oper. Res.*, **58** (2024), 2367–2378. https://doi.org/10.1051/ro/2024098
7. C. T. Hoàng, Perfect graphs, PhD thesis, McGill University, Montreal, 1985.
8. V. Giakoumakis, F. Roussel, H. Thuillier, On P_4-tidy graphs, *Discrete Math. Theor.*, **1** (1997), 17–41. https://doi.org/10.46298/dmtcs.232
9. B. Brešar, T. Kos, G. Nasini, P. Torres, Total dominating sequences in trees, split graphs, and under modular decomposition, *Discrete Optim.*, **28** (2018), 16–30. https://doi.org/10.1016/j.disopt.2017.10.002
10. P. Torres, Grundy double domination number: bounds, graph operations, and efficient computation for P_4-tidy graphs, *Bull. Malays. Math. Sci. Soc.*, **49** (2026), 52. https://doi.org/10.1007/s40840-026-02055-z
11. S. Klein, C. P. de Mello, A. Morgana, Recognizing well covered graphs of families with special P_4-components, *Graph. Combinator.*, **29** (2013), 553–567. https://doi.org/10.1007/s00373-011-1123-1
12. M. R. Cappelle, E. Coelho, L. R. Foulds, H. J. Longo, Open-independent, open-locating-dominating sets: structural aspects of some classes of graphs, *Discrete Math. Theor.*, **24** (2022), #5. https://doi.org/10.46298/dmtcs.8440
13. L. C. D. Carvalho, M. R. Cappelle, k-Independência em grafos P_4-esparsos, *Anais do LIV Simpósio Brasileiro de Pesquisa Operacional*, **54** (2022), 152990.
14. B. Courcelle, J. A. Makowsky, U. Rotics, Linear time solvable optimization problems on graphs of bounded clique-width, *Theor. Comput. Syst.*, **33** (2000), 125–150. https://doi.org/10.1007/s002249910009
15. N. Betzler, R. Bredereck, R. Niedermeier, J. Uhlmann, On bounded-degree vertex deletion parameterized by treewidth, *Discrete Appl. Math.*, **160** (2012), 53–60. https://doi.org/10.1016/j.dam.2011.08.013
16. M. U. Gerber, D. Kobler, Algorithms for vertex-partitioning problems on graphs with fixed clique-width, *Theor. Comput. Sci.*, **299** (2003), 719–734. https://doi.org/10.1016/S0304-3975(02)00725-9
17. M. Lampis, Parameterized approximation schemes using graph widths, In: *Automata, languages, and programming (ICALP 2014)*, Berlin: Springer, 2014, 775–786. https://doi.org/10.1007/978-3-662-43948-7_64
18. B. Jamison, S. Olariu, A tree representation for P_4-sparse graphs, *Discrete Appl. Math.*, **35** (1992), 115–129. https://doi.org/10.1016/0166-218X(92)90036-A
19. M. R. Cappelle, L. C. D. Carvalho, Computational verification for k-independence in P_4-tidy graphs (Version v1.0.1) [Computer software], Zenodo, 2026. https://doi.org/10.5281/zenodo.22864016
20. R. M. McConnell, J. P. Spinrad, Modular decomposition and transitive orientation, *Discrete Math.*, **201** (1999), 189–241. https://doi.org/10.1016/S0012-365X(98)00319-7
21. M. Tedder, D. Corneil, M. Habib, C. Paul, Simpler linear-time modular decomposition via recursive factorizing permutations, In: *Automata, languages and programming (ICALP 2008)*, Berlin: Springer, 2008, 634–645. https://doi.org/10.1007/978-3-540-70575-8_52

22. B. Courcelle, J. A. Makowsky, U. Rotics, On the fixed parameter complexity of graph enumeration problems definable in monadic second-order logic, *Discrete Appl. Math.*, **108** (2001), 23–52. https://doi.org/10.1016/S0166-218X(00)00221-3



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)