Searching for an intruder on graphs and their subdivisions

Anton Bernshteyn*

School of Mathematics
Georgia Institute of Technology
Atlanta, GA, USA

bahtoh@gatech.edu

Eugene Lee

Department of Mathematical Sciences
Carnegie Mellon University
Pittsburgh, PA, USA

eleehuaj@andrew.cmu.edu

Abstract

In this paper we analyze a variant of the pursuit-evasion game on a graph G where the intruder occupies a vertex, is allowed to move to adjacent vertices or remain in place, and is 'invisible' to the searcher, meaning that the searcher operates with no knowledge of the position of the intruder. On each stage, the searcher is allowed to inspect an arbitrary set of k vertices. The minimum k for which the searcher can guarantee the capture of the intruder is called the *inspection number* of G. We also introduce and study the *topological inspection number*, a quantity that captures the limiting behavior of the inspection number under subdivisions of G. Our central theorem provides a full classification of graphs with topological inspection number up to 3.

Mathematics Subject Classifications: 05C57, 05C75

1 Introduction

All graphs in this paper are finite, undirected, and simple. A class of graph-theoretic problems that has been the topic of much interest is that of pursuit-evasion games, where the object is to capture an intruder who is allowed to move within a graph in some manner. The general nature of this problem lends itself to many variants; whether the intruder occupies vertices or edges of the graph, the amount of information revealed to the searchers as well as the constraints on their movement are all facets of the problem that can be modified to ask different questions. The version we analyze in this paper is the following:

Definition 1 (Zero-Visibility Search Game)**.** For a positive integer k, the **zero-visibility k-search game** is a two-player game played on a graph G, where one player serves as the **searcher**, while the other plays as the **intruder**. Both players have full knowledge of the graph.

The intruder starts the game by occupying a vertex of their choice. At all times, the location of the intruder is unknown to the searcher (in other words, the intruder is 'invisible'). The players alternate turns, beginning with the searcher. On each turn, the searcher 'inspects' an arbitrary k-element set of vertices, and if the intruder is presently at any of those vertices, the game ends and the searcher wins. Otherwise, on the intruder's turn, they can choose to move to an adjacent vertex or remain at the same vertex. We emphasize that the searcher does not know if and where the intruder moves.

A winning strategy for the searcher is a finite sequence of moves which guarantees that the intruder will always be caught, regardless of the intruder's starting location and moves. The **inspection number** of the graph, denoted $\mathsf{in}(G)$, is the minimum k such that a winning strategy for the searcher exists.

This problem has a natural equivalent formulation in terms of curbing the spread of an infection on the vertices of G. Suppose that each vertex can be in one of two states: *cleared* or *contaminated*. Initially, all vertices are contaminated. On every move, we may clear any k vertices. After this, every cleared vertex that has a contaminated neighbor becomes contaminated again. The inspection number of G is then equal to the smallest k for which all vertices of G can be cleared in finitely many moves (see Proposition 10).

Very similar pursuit-evasion games have been considered previously. Tošić [16] introduced the *zero-visibility Cops & Robber game*, which differs from ours in that instead of examining arbitrary sets of k vertices, the searcher controls k tokens that occupy vertices of G and can only be moved along edges. This model was also studied by Berger, Gilbers, Grüne, and Klein [4] and Dereniowski, Dyer, Tifenbach, and Yang [9]. Another related model was studied in [3] by Aydinian, Cicalese, Deppe, and Lebedev. In their model, the searcher can only learn whether the intruder is located at one of the examined vertices (but not at which vertex specifically). The game investigated in [14] by Haslegrave, in [7] by Britnell and Wildon, and in [1] by Abramovskaya, Fomin, Golovach, and Pilipczuk, called in the latter the *Hunters & Rabbit game*, is almost the same as ours except that the intruder is required to move to an adjacent vertex on every move. For further references, see the surveys [2] by Alspach, [13] by Fomin and Thilikos, and [6] by Bonato and Yang.

It is easy to see that the complete graph K_n has inspection number n, paths have inspection number 2, and cycles have inspection number 3. We also compute the inspection number of rectangular grid graphs:

Theorem 2 (See Theorem 14)**.** *For each $n, m \geqslant 2$, the inspection number of the $n \times m$ grid graph is $\min\{n, m\} + 1$.*

Trees can have arbitrarily high inspection numbers:

Theorem 3 (See Theorem 15)**.** *For every k, there is a tree T of maximum degree 3 with $\mathsf{in}(T) > k$.*
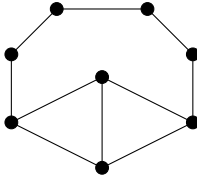
Figure 1: This graph (a subdivision of K_4) has pathwidth 3 and hence monotonic inspection number 4, but its ordinary inspection number is 3 (see Example 22).

Our proof of Theorem 3 is surprisingly intricate. In general, proving lower bounds on the inspection number appears challenging.

An interesting aspect of pursuit-evasion games is whether or not they admit monotonic optimal strategies. Here a strategy is *monotonic* if cleared vertices never become contaminated again. Some pursuit-evasion games admit optimal monotonic strategies [5, 15], while some do not [17]. In their investigation of the zero-visibility Cops & Robber game, Dereniowski et al. [9] established sharp bounds on the number of cops needed for a monotonic winning strategy in terms of the pathwidth of the underlying graph. We prove the following result in the same vein:

Theorem 4 (See Theorem 20). *For every connected graph G, the minimum k such that the searcher has a monotonic winning strategy in the zero-visibility k-search game on G (the* **monotonic inspection number**) *is equal to* $\mathsf{pw}(G) + 1$, *where* $\mathsf{pw}(G)$ *denotes the pathwidth of G.*

Theorem 20 implies that $\mathsf{in}(G) \leqslant \mathsf{pw}(G) + 1$ for all connected graphs G. The $n \times m$ grid graph meets this bound with equality, since it has pathwidth $\min\{n, m\}$ [12] (note that the proof in the cited paper is via considering another pursuit-evasion game) and inspection number $\min\{n, m\} + 1$ by Theorem 2. On the other hand, there are graphs G with $\mathsf{in}(G) < \mathsf{pw}(G) + 1$ (and hence without a monotonic optimal search strategy). One such example is shown in Fig. 1. Our further results imply that there are graphs with inspection number 3 and arbitrarily high pathwidth (see Theorem 6).

Another interesting direction is to study pursuit-evasion games on subdivisions of graphs, see, e.g., [8]. Edge subdivisions can both increase and decrease the inspection number. For example, the complete graph K_4 has inspection number 4 (see Example 11). The subdivision of K_4 shown in Fig. 1 has inspection number 3, but it also has a further subdivision with inspection number 4 (see Theorem 32). We hence wish to investigate the "limiting" behavior of the inspection number under edge subdivisions, captured precisely in the following notion:

Definition 5. The **topological inspection number** of a graph G, denoted $\mathsf{in}_t(G)$, is the minimum k such that for every subdivision H of G, there is a further subdivision H' of H with $\mathsf{in}(H') \leqslant k$.

Equivalently, $\mathsf{in}_t(G)$ is the minimum k such that for all ℓ, G has a subdivision H with $\mathsf{in}(H) \leqslant k$ in which every edge of G is subdivided at least ℓ times.
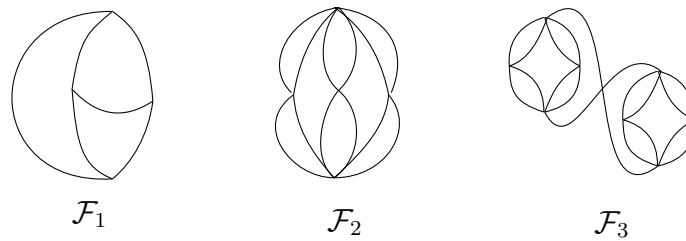
Figure 2: A representative example of each class of forbidden subgraphs. The family $\mathcal{F}_1$ comprises the subdivisions of K_4, while the other two families are more complicated (see Definition 26).

The topological inspection number can deviate significantly from the inspection number. For instance, by Theorem 15, trees can have arbitrarily high inspection numbers, but we shall prove the following:

Theorem 6. *The topological inspection number of every tree is at most* 3.

Theorem 6 also implies that there are graphs with inspection number 3 and arbitrarily high pathwidth, since the pathwidth does not decrease under edge subdivisions and there are trees of arbitrarily high pathwidth.

Theorem 6 is a corollary of our main result, which provides a complete characterization of graph with topological inspection number at most 3. It is not hard to see that the only connected graphs with topological inspection number at most 2 are paths (see Proposition 24). The case of graphs with topological inspection number at most 3 turns out to be significantly more involved:

Theorem 7 (See Theorem 30). *Let G be a connected graph with $|V(G)| \geqslant 2$. The following are equivalent:*

  (1) *G has topological inspection number at most* 3.

  (2) *G does not contain as a subgraph any element of $\mathcal{F}_1, \mathcal{F}_2$, or $\mathcal{F}_3$ (see Definition 26).*

  (3) *G has a simple generalized series-parallel decomposition (see Definitions 27–29).*

This equivalence gives two different ways of detecting if a graph has topological inspection number at most 3. Item (2) describes such graphs by means of an explicit list of forbidden subgraphs. The forbidden subgraphs form three infinite families. The first family, $\mathcal{F}_1$, comprises the subdivisions of K_4. The other two families, $\mathcal{F}_2$ and $\mathcal{F}_3$, are described in Definition 26. Representative examples of each family of forbidden subgraphs are depicted in Fig. 2. All the graphs in these families contain cycles, so the equivalence (1) $\iff$ (2) yields Theorem 6.
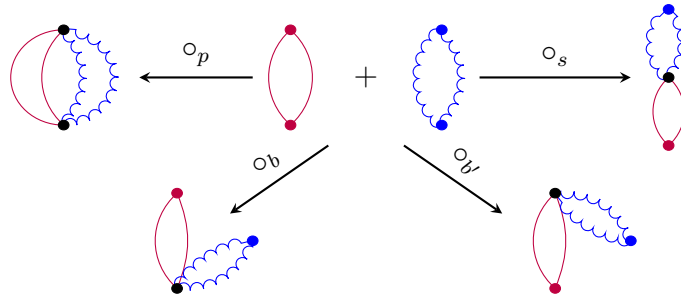
Figure 3: Pictorial representations of the four permissible operations defining the class of generalized series-parallel graphs. A decomposition is *simple* if each of the $\circ_p$ operations in it has one of the subgraphs involved containing an edge separating its *terminals*, the two distinguished vertices.

Item (3) gives a constructive characterization of graphs with topological inspection number at most 3. Roughly speaking, it says that a graph with topological inspection number at most 3 can be inductively built up from individual edges via certain permissible operations. The four permissible operations are the *series operation* $\circ_s$, the *parallel operation* $\circ_p$, and the two *branch operations* $\circ_b$ and $\circ_{b'}$. These operations are illustrated in Fig. 3. Graphs that can be constructed from individual edges using these operations are called *generalized series-parallel* (*GSP* for short); see Definition 27 for details.

It is well-known that 2-connected K_4-subdivision-free graphs are *series-parallel*, i.e., they can be built from individual edges using only the operations $\circ_s$ and $\circ_p$ [11, Theorem 2]. In general, connected K_4-subdivision-free graphs are GSP (see §8.1). Since subdivisions of K_4 belong to the forbidden family $\mathcal{F}_1$, item (2) of Theorem 7 implies that every connected graph G with $\mathsf{in}_t(G) \leqslant 3$ must be GSP. However, as the graphs in $\mathcal{F}_2$ and $\mathcal{F}_3$ demonstrate, not all GSP graphs have topological inspection number at most 3. The main point of (3) is that a graph G with $\mathsf{in}_t(G) \leqslant 3$ must admit a GSP decomposition that is *simple*, which means that every time the parallel operation is used, at least one of the graphs it is applied to has to contain an edge separating the two distinguished vertices. The details are given in Definition 29.

Our proof of the equivalence (1) $\iff$ (3) is constructive; that is, given a graph G with $\mathsf{in}_t(G) \leqslant 3$, it explicitly describes a subdivision of G with a successful search strategy.

Our characterization of graphs with topological inspection number at most 3 via forbidden subgraphs shows that every such graph must be planar, since every non-planar graph contains a subdivision of K_4. Furthermore, some forbidden graphs in $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$ are planar. On the other hand, none of the forbidden subgraphs are outerplanar, so every outerplanar graph has topological inspection number at most 3. Also, there exist graphs which are not outerplanar but have topological inspection number 3 ($K_{2,3}$ is such an example). Thus, graphs with topological inspection number at most 3 form a natural class strictly between outerplanar and planar graphs.

The rest of the paper is organized as follows.

- Section 2 introduces the notation to be used in the rest of the paper.

- Section 3 computes bounds on the inspection number for several special classes of graphs. In particular, it contains the proofs of Theorems 2 and 15.

- Section 4 investigates monotonic search strategies and proves Theorem 20.

- Section 5 introduces the topological inspection number and classifies graphs with topological inspection number at most 2.

- Section 6 introduces the necessary definitions for the statement of Theorem 7.

- Section 7 proves that each graph in $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$ has topological inspection number at least 4 (this is implication (1) $\Longrightarrow$ (2) of Theorem 7).

- Section 8 proves that every graph without any forbidden subgraph has a simple generalized series-parallel decomposition (this is implication (2) $\Longrightarrow$ (3) of Theorem 7).

- Section 9 proves that every graph with a simple generalized series-parallel decomposition has topological inspection number at most 3 (this is implication (3) $\Longrightarrow$ (1) of Theorem 7).

- Section 10 discusses several future directions of study, as well as some open problems.

Acknowledgments

2 Preliminaries

2.1 Graph-theoretic notation

Given a graph G, we denote by $V(G)$ and $E(G)$ its vertex and edge sets, respectively. If $S \subseteq V(G)$, $N_G(S)$ denotes the set of neighbors of S in G. We use $G[S]$ to denote the subgraph of G induced by S, and $G - S$ is the subgraph induced by $V(G) \setminus S$. We call a set S of vertices G-**connected** if $G[S]$ is connected. For $v \in V(G)$ and $r \in \mathbb{N}$, $B_G(v, r)$ denotes the **ball** of radius r around v in G, i.e., the set of all the vertices of G at distance at most r from v. The **boundary** of S, denoted $\partial_G(S)$, is the set of all the vertices in S that have a neighbor in $V(G) \setminus S$. Note that if H is a subgraph of G and $S \subseteq V(H)$, then $\partial_H(S) \subseteq \partial_G(S)$.

We say a set $S \subseteq V(G)$ **separates** a collection of nonempty sets of vertices $\mathcal{F} \subseteq \mathcal{P}(V(G) \setminus S)$ if every connected component of $G - S$ has nonempty intersection with at most one member of $\mathcal{F}$. We naturally extend this notion to separating vertices by identifying vertices with singleton sets.

Lemma 8. *If G is k-connected and $S \subseteq V(G)$ satisfies $|S| \geqslant k$, then $|\partial_G(S)| \geqslant k$ or $S = V(G)$.*

Proof. Note that $\partial_G(S)$ separates $S \setminus \partial_G(S)$ from $V(G) \setminus S$. If $|\partial_G(S)| < |S| < V(G)$, then these two sets are nonempty, and hence $|\partial_G(S)| \geqslant k$. Otherwise, we have $S = V(G)$ or $|\partial_G(S)| = |S| \geqslant k$. ∎

2.2 Searches

Let us now introduce some notation useful for describing search strategies.

Definition 9 (Searches). A k-**search** of length ℓ on a graph G is a sequence $\mathcal{S} = \{S_t \subseteq V(G)\}_{t=1}^{\ell}$ for some $\ell \in \mathbb{N}$ where $|S_t| \leqslant k$ for each t. Note the convention of denoting searches by calligraphic letters and the individual sets by the same non-calligraphic letter with a subscript. We refer to S_t (in the context of a search $\mathcal{S}$) as the **set of vertices searched at turn** t. We write $|\mathcal{S}| = \ell$ to denote the length of $\mathcal{S}$. We sometimes refer to k as the **search size**. We will occasionally refer to k-searches as **searches** when the specific search size is not important.

Given an **initial set of cleared vertices** $A \subseteq V(G)$ (which is usually taken to be $\varnothing$), each k-search $\mathcal{S}$ of length ℓ has two associated sequences of sets of vertices $\{\mathsf{PC}_t(\mathcal{S}, A)\}_{t=1}^{\ell}$ and $\{\mathsf{FC}_t(\mathcal{S}, A)\}_{t=0}^{\ell}$, the **pre-cleared** and **fully cleared** vertices respectively after each turn. These are defined recursively by:

$$\mathsf{FC}_0(\mathcal{S}, A) = A;$$
$$\mathsf{PC}_{t+1}(\mathcal{S}, A) = \mathsf{FC}_t(\mathcal{S}, A) \cup S_{t+1};$$
$$\mathsf{FC}_{t+1}(\mathcal{S}, A) = \mathsf{PC}_{t+1}(\mathcal{S}, A) \setminus \partial_G(\mathsf{PC}_{t+1}(\mathcal{S}, A)).$$

When $A = \varnothing$, we will typically omit it from the notation, thus writing $\mathsf{PC}_t(\mathcal{S})$ and $\mathsf{FC}_t(\mathcal{S})$ for $\mathsf{PC}_t(\mathcal{S}, \varnothing)$ and $\mathsf{FC}_t(\mathcal{S}, \varnothing)$ respectively. We use $\mathsf{FC}_*(\mathcal{S}, A)$ (resp. $\mathsf{PC}_*(\mathcal{S}, A)$) to denote the set $\mathsf{FC}_\ell(\mathcal{S}, A)$ (resp. $\mathsf{PC}_\ell(\mathcal{S}, A)$) of fully cleared (resp. pre-cleared) vertices after the final turn.

We say a k-search $\mathcal{S}$ is **successful** if $\mathsf{FC}_*(\mathcal{S}) = V(G)$.

Proposition 10. *The inspection number of a graph G (as given by Definition 1) is equal to the minimum k such that there exists a successful k-search on G.*

Proof. We can consider the natural correspondence between strategies for the searcher and k-searches by identifying the sets of vertices inspected at turn i with the ith term of the k-search. A simple inductive argument shows that given this correspondence, the set of all the possible locations of the intruder after the searcher's ith turn is precisely $V(G) \setminus \mathsf{PC}_i(\mathcal{S})$ and the set of all the possible locations of the intruder after the intruder's ith turn is precisely $V(G) \setminus \mathsf{FC}_i(\mathcal{S})$. The searcher's strategy is winning if and only if the set of possible locations is reduced to $\varnothing$ by a finite sequence of moves, which is equivalent to the search being successful. ∎

It is clear that for any subgraph H of G, we have $\mathsf{in}(H) \leqslant \mathsf{in}(G)$, as any successful search on G, with the vertex sets appropriately restricted, is also a successful search on H.

Example 11. $\mathsf{in}(K_n) = n$.

Proof. Suppose $\mathcal{S}$ is a k-search on K_n for some $k < n$. Note that if $\mathsf{FC}_i(\mathcal{S}) = \varnothing$, then $\mathsf{PC}_{i+1}(\mathcal{S}) \neq V(K_n)$, so $\mathsf{FC}_{i+1}(\mathcal{S}) = \varnothing$ as well. Hence $\mathcal{S}$ cannot be successful. On the other hand, checking all the vertices on the first move gives a successful n-search, and hence $\mathsf{in}(K_n) = n$. ∎

3 Some computations

3.1 Sets with small boundary

Our primary tool for proving lower bounds on the inspection number is the following proposition:

Proposition 12. *Let G be a connected graph. If $\mathsf{in}(G) \leqslant k$, then for each $1 \leqslant i \leqslant |V(G)|$, there exists some $C \subseteq V(G)$ such that $|\partial_G(C)| < k$ and $i - k < |C| < i$.*

Proof. Let $\mathcal{S}$ be a successful k-search on G. For each $1 \leqslant i \leqslant |V(G)|$, consider the minimal t such that $|\mathsf{PC}_t(\mathcal{S})| \geqslant i$. We claim that $C = \mathsf{PC}_{t-1}(\mathcal{S})$ works. Indeed, $|\mathsf{PC}_{t-1}(\mathcal{S})| < i$, so we can write

$$|\mathsf{PC}_{t-1}(\mathcal{S})| < i \leqslant |\mathsf{PC}_t(\mathcal{S})| \leqslant |\mathsf{FC}_{t-1}(\mathcal{S})| + k = |\mathsf{PC}_{t-1}(\mathcal{S})| - |\partial_G(\mathsf{PC}_{t-1}(\mathcal{S}))| + k.$$

Thus, we have $|\partial_G(\mathsf{PC}_{t-1}(\mathcal{S}))| < k$. Furthermore,

$$i - k \leqslant |\mathsf{PC}_{t-1}(\mathcal{S})| - |\partial_G(\mathsf{PC}_{t-1}(\mathcal{S}))| < |\mathsf{PC}_{t-1}(\mathcal{S})| < i,$$

where we use that $\partial_G(\mathsf{PC}_{t-1}(\mathcal{S}))$ is nonempty. ∎

3.2 Grid graphs

Here we compute the inspection number of rectangular grid graphs.

Definition 13 (Grid Graphs). An $n \times m$ **grid graph**, denoted $\Gamma_{n,m}$, comprises the vertices

$$\{(i,j) \in \mathbb{Z}^2 \mid 0 \leqslant i < n, 0 \leqslant j < m\},$$

with vertices (a,b) and (c,d) adjacent if and only if $|a - c| + |b - d| = 1$.

Visually the first coordinate indexes the row of a vertex, and the second one the column. Rows increase in index from bottom to top, and columns increase in index from left to right.

Theorem 14. *For each $n, m \geqslant 2$, $\mathsf{in}(\Gamma_{n,m}) = \min\{n, m\} + 1$.*

Proof. Let $G = \Gamma_{n,m}$. Suppose, without loss of generality, that $n \geqslant m$, so our goal is thus to show that $\mathsf{in}(G) = m + 1$. For convenience, we relabel the vertices of G, writing (a, b) as v_{am+b} (see Fig. 4).

To prove that $\mathsf{in}(G) \leqslant m + 1$, consider the $(m + 1)$-search $\mathcal{S}$ defined via

$$S_t := \{v_{t+\ell} \mid 0 \leqslant \ell \leqslant m\} \quad \text{for all } 1 \leqslant t \leqslant m(n - 1).$$
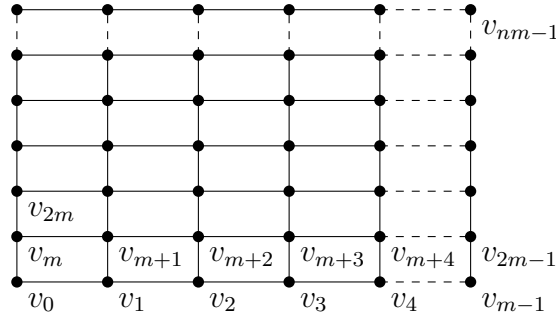
Figure 4: The graph $\Gamma_{n,m}$.

Note that for each v_i, its neighbors v_j satisfy $j \leqslant i + m$. If for some t, $\{v_i \mid i \leqslant t\} \subseteq \mathsf{FC}_t(\mathcal{S})$, then we have $\{v_i \mid i \leqslant t + m + 1\} \subseteq \mathsf{PC}_{t+1}(\mathcal{S})$, and hence $\{v_i \mid i \leqslant t + 1\} \subseteq \mathsf{FC}_{t+1}(\mathcal{S})$. Induction thus shows that the search is successful, so $\mathsf{in}(G) \leqslant m + 1$, as desired.

We now wish to show that no successful m-search exists. We shall analyze the sets with boundaries of size less than m and then use Proposition 12 to conclude the argument.

Claim 14.a. *Any set* $S \subseteq V(G)$ *with* $|\partial_G(S)| \leqslant m - 1$ *satisfies*

$$|S| \leqslant \frac{m(m-1)}{2} \quad or \quad |S| \geqslant mn - \frac{(m-2)(m-1)}{2}.$$

$\triangleright$ Denote the ith column and jth row of G by c_i and r_j respectively. For brevity, let a row or column be **empty/full** if it is disjoint from/contained in S respectively. Note that each row or column is either empty, full, or contains a boundary vertex of S.

If there exist both an empty column and a full column, then every row is neither empty nor full and hence there are at least n boundary vertices of S, a contradiction. Similarly, if there is neither an empty column nor a full column, then every column contains a boundary vertex of S and so there are at least m boundary vertices of S, a contradiction.

Now suppose there exists an empty column c_k but no full columns. Take some other arbitrary column c_i and consider the segments of each row between c_i and c_k. For $|c_i \cap S|$ of them, they contain a vertex in c_i, and none of them contain a vertex in c_k, so they contain at least $|c_i \cap S|$ boundary vertices of S. Furthermore, every column not contained between c_i and c_k inclusive is disjoint from each of these segments. If $m(i)$ is the number of nonempty columns not contained between c_i and c_k inclusive, then $|\partial_G(S)| \geqslant |c_i \cap S| + m(i)$, as such columns are neither empty nor full and hence contain at least one boundary vertex of S each. Thus,

$$m - 1 - m(i) \geqslant |c_i \cap S|. \tag{$*$}$$

Now let $I = \{i \mid c_i \text{ is nonempty}\}$. Then

$$\sum_{i \in I} m(i) \geqslant \binom{|I|}{2},$$

as for each pair of distinct columns c_i, c_j, either c_i is not between c_j and c_k or c_j is not between c_i and c_k. Summing ($*$) over I gives

$$|S| = \sum_{i \in I} |c_i \cap S| \leqslant \sum_{i \in I} (m - 1 - m(i))$$

$$\leqslant |I|(m - 1) - \binom{|I|}{2} = \frac{|I|(2m - |I| - 1)}{2} \leqslant \frac{(m - 1)m}{2}.$$

The remaining case where there exists a full column c_k is analogous, except we get the boundary vertices on the segments for which c_i does not contain a vertex of S, and thus the boundary vertex in c_i is distinct from the boundary vertices of each segment. The analog of ($*$) is hence

$$|\partial_G(S)| \geqslant |c_i \setminus S| + 1 + m'(i) \quad \implies \quad m - 2 - m'(i) \geqslant |c_i \setminus S|,$$

where $m'(i)$ is the number of non-full columns not contained between c_i and c_k. Via the same bound for the sum of $m'(i)$, we get

$$|V(G) \setminus S| \leqslant \sum_{i \in J} (m - 2 - m'(i)) \leqslant \frac{|J|(2m - |J| - 3)}{2} \leqslant \frac{(m - 2)(m - 1)}{2},$$

where $J = \{i \mid c_i \text{ is not full}\}$. Therefore, $|S| \geqslant mn - \frac{(m-2)(m-1)}{2}$, as desired. ◀

Claim 14.a implies that there does not exist a set S with $|\partial_G(S)| < m$ and

$$\frac{m(m-1)}{2} < |S| < \frac{m(m-1)}{2} + m = m^2 - \frac{(m-2)(m-1)}{2} \leqslant mn - \frac{(m-2)(m-1)}{2},$$

contradicting Proposition 12 and completing the proof. ∎

3.3 Trees

In this subsection we show that trees can have arbitrarily large inspection number.

Theorem 15. *For every k, there is a tree T of maximum degree 3 with $\mathsf{in}(T) > k$.*

Proof. Take any $N > \log_2(k) + 2$ and consider the perfect binary tree T of depth $L = 20kN$, i.e., a rooted tree in which every non-leaf vertex has exactly two children and the distance from the root to every leaf is equal to L. We wish to compute the possible sizes of subsets of $V(T)$ with fewer than k boundary vertices and then apply Proposition 12 to conclude that $\mathsf{in}(T) > k$.

Consider an arbitrary subset $S \subseteq V(T)$ with $|\partial_T(S)| < k$. For each $0 \leqslant i \leqslant L$, let S_i be the set of the vertices of S at depth $L - i$ (so S_0 comprises the leaf vertices in S, while S_L will be the root if it is in S). We define two subsets $A_i, B_i \subseteq S_i \cap \partial_T(S)$ by

$$A_i = \{x \in S_i \cap \partial_T(S) \mid \text{the parent of } x \text{ is not in } S\};$$
$$B_i = \{x \in S_i \cap \partial_T(S) \mid \text{not every child of } x \text{ is in } S\}.$$

This gives us a recursive bound on the size of S_i for $i < L$:

$$2\,|S_{i+1}| + |A_i| - 2\,|B_{i+1}| \leqslant |S_i| \leqslant 2\,|S_{i+1}| + |A_i| - |B_{i+1}|,$$

noting that every vertex in B_{i+1} has either 0 or 1 child in S. Iterating this bound gives the following relation between the sizes of any two S_is:

$$2^\ell\,|S_{i+\ell}| + \sum_{j=0}^{\ell-1} 2^j\left(|A_{i+j}| - 2\,|B_{i+j+1}|\right) \leqslant |S_i| \leqslant 2^\ell\,|S_{i+\ell}| + \sum_{j=0}^{\ell-1} 2^j\left(|A_{i+j}| - |B_{i+j+1}|\right). \quad (\spadesuit)$$

We denote by $R(i;\ell)$ the 'slice' $\bigcup_{j=0}^{\ell} S_{i+j}$. Summing $(\spadesuit)$ gives:

$$\left(2^{\ell+1} - 1\right)|S_\ell| + \sum_{j=0}^{\ell-1}\left(2^{j+1} - 1\right)\left(|A_j| - 2\,|B_{j+1}|\right) \leqslant |R(0;\ell)|$$

$$\leqslant \left(2^{\ell+1} - 1\right)|S_\ell| + \sum_{j=0}^{\ell-1}\left(2^{j+1} - 1\right)\left(|A_j| - |B_{j+1}|\right). \quad (\heartsuit)$$

By Proposition 12, if $\mathsf{in}(T) \leqslant k$, then there exists a set $S \subseteq V(T)$ with $|\partial_T(S)| < k$ such that

$$\left||S| - \sum_{j=1}^{4k+1} 2^{L-jN}\right| < k. \quad (\diamondsuit)$$

For the remainder of the proof, we fix any such set S.

Claim 15.a. *For each $1 \leqslant a \leqslant 2k$, there is $1 \leqslant j \leqslant 2N$ such that $|A_{L-2aN-j}|$ or $|B_{L-2aN-j}|$ is nonzero.*

▷   Note that $|S| = |R(0; L-2aN-1)| + |R(L-2aN; 2aN)|$. Since $R(L-2aN; 2aN)$ is a subset of a perfect binary tree of depth $2aN$,

$$|R(L-2aN; 2aN)| \leqslant 2^{2aN+1} - 1 \leqslant 2^{5kN}.$$

Suppose for the sake of contradiction that

$$|A_{L-2aN-j}| = |B_{L-2aN-j}| = 0 \quad \text{for all } 1 \leqslant j \leqslant 2N. \quad (\clubsuit)$$

Let $K = \left(2^{L-2aN} - 1\right)|S_{L-2aN-1}|$. Then, by $(\heartsuit)$, we have

$$|R(0; L-2aN-1)| \geqslant K + \sum_{j=0}^{L-2aN-2}\left(2^{j+1} - 1\right)\left(|A_j| - 2\,|B_{j+1}|\right)$$

$$[\text{by } (\clubsuit)] \qquad \geqslant K - \sum_{j=0}^{L-2(a+1)N-2}\left(2^{j+1} - 1\right)2\,|B_{j+1}|$$

$$[\text{since } \sum|B_i| \leqslant |\partial_T(S)| < k] \qquad \geqslant K - 2^{L-2(a+1)N}k.$$

Similarly,

$$|R(0; L - 2aN - 1)| \leqslant K + \sum_{j=0}^{L-2aN-2} \left(2^{j+1} - 1\right) \left(|A_j| - |B_{j+1}|\right)$$

$$[\text{by } (\clubsuit)] \qquad \leqslant K + \sum_{j=0}^{L-2(a+1)N-1} \left(2^{j+1} - 1\right) |A_j|$$

$$[\text{since } \textstyle\sum |A_i| \leqslant |\partial_T(S)| < k] \qquad \leqslant K + 2^{L-2(a+1)N} k.$$

Since $|S_{L-2aN-1}| \leqslant 2^{2aN+1} \leqslant 2^{5kN}$, K is within 2^{5kN} of a multiple of 2^{L-2aN}. The maximum difference that $|S|$ could attain from a multiple of 2^{L-2aN} is therefore

$$|R(L - 2aN; 2aN)| + |S_{L-2aN-1}| + 2^{L-2(a+1)N}k \leqslant 2^{5kN} + 2^{5kN} + 2^{L-2(a+1)N}k$$

$$< 2^{L-2(a+1)N + \log_2(k) + 1}.$$

On the other hand, by $(\diamondsuit)$, $|S|$ differs from a multiple of 2^{L-2aN} at least by

$$2^{L-(2a+1)N} - k > 2^{L-(2a+1)N-1},$$

which is a contradiction since $N > \log_2(k) + 2$. ◄

Applying Claim 15.a to each $1 \leqslant a \leqslant 2k$ yields $2k$ distinct values i such that $|A_i|$ or $|B_i|$ is nonzero, and therefore $\sum |A_i| + \sum |B_i| \geqslant 2k$. This is a contradiction since $\sum |A_i|$, $\sum |B_i| \leqslant |\partial_T(S)| < k$. ∎

4 Monotonic searches

Definition 16 (Monotonic Searches). A search $\mathcal{S}$ is called **monotonic** if $\mathsf{FC}_i(\mathcal{S}) \subseteq \mathsf{FC}_j(\mathcal{S})$ for every $i \leqslant j$. Let $\mathsf{in}_m(G)$ be the minimum k for which there exists a successful monotonic k-search on G.

Let us also recall the standard definition of the pathwidth of a graph.

Definition 17 (Pathwidth). A **path decomposition** of a graph G is a sequence of subsets $\{X_i \subseteq V(G)\}_{i=1}^{\ell}$ fulfilling the following two criteria:

- For each edge $\{a, b\} \in E(G)$, there exists some i such that $\{a, b\} \subseteq X_i$.

- For each $1 \leqslant i < j < k \leqslant \ell$ and $v \in V(G)$, if $v \in X_i$ and $v \in X_k$, then $v \in X_j$.

The **width** of such a path decomposition is $\max_{1 \leqslant i \leqslant \ell} |X_i| - 1$ (if the decomposition is empty we set the width to be 0). The **pathwidth** of a graph, denoted $\mathsf{pw}(G)$, is the minimum width among all its path decompositions.

Proposition 18. *For any connected graph G, $\mathsf{in}_m(G) \leqslant \mathsf{pw}(G) + 1$.*

Proof. The statement is trivial if G has only one vertex, so we may assume that G does not have isolated vertices, and thus every vertex must be included in at least one set in each path decomposition of G.

Let $\mathcal{S} = \{S_i \subseteq V(G)\}_{i=1}^{k}$ be a path-decomposition of G of width $\mathsf{pw}(G)$, i.e., $|S_i| \leqslant \mathsf{pw}(G) + 1$ for each i, for each edge of G there exists i such that S_i contains both its endpoints, and for all $i \leqslant j \leqslant k$, $S_i \cap S_k \subseteq S_j$. Note that for each $v \in V(G)$, the set of indices i for which $v \in S_i$ forms a consecutive sequence of integers. We claim that $\mathcal{S}$ is a successful monotonic $(\mathsf{pw}(G) + 1)$-search of G.

First, we show by induction that for each t, if $v \notin S_i$ for all $i > t$, then $v \in \mathsf{FC}_t(\mathcal{S})$. Suppose the claim holds for all $t < t_0$, and suppose $v \notin S_i$ for every $i > t_0$. Consider any $u \in N_G(v)$, and let t_u be maximal such that $u \in S_{t_u}$. If $t_u \geqslant t_0$, then $u \in S_{t_0}$ as $\{u, v\} \subseteq S_i$ for some $i \leqslant t_0$, and so $u \in \mathsf{PC}_{t_0}(\mathcal{S})$. Otherwise, $u \notin S_i$ for each $i > t_0 - 1$, so by the induction hypothesis $u \in \mathsf{FC}_{t_0-1}(\mathcal{S}) \subseteq \mathsf{PC}_{t_0}(\mathcal{S})$. A similar argument gives that $v \in \mathsf{PC}_{t_0}(\mathcal{S})$. Hence v and each of its neighbors are in $\mathsf{PC}_{t_0}(\mathcal{S})$, so $v \in \mathsf{FC}_{t_0}(\mathcal{S})$, completing the induction. This shows that the search is successful.

Now we show that the search is monotonic. Suppose that $v \in \mathsf{FC}_t(\mathcal{S}) \setminus \mathsf{FC}_{t+1}(\mathcal{S})$. Then there exists some $u \in N_G(v)$ such that $u \in \mathsf{PC}_t(\mathcal{S}) \setminus \mathsf{PC}_{t+1}(\mathcal{S})$. Since $u \in \mathsf{PC}_t(\mathcal{S})$, there is some $t_0 \leqslant t$ such that $u \in S_{t_0}$. On the other hand, since $u \notin \mathsf{PC}_{t+1}(\mathcal{S})$, we have $u \notin S_{t+1}$. This implies that $u \notin S_{t_1}$ for all $t_1 \geqslant t + 1$, and hence $u \notin \mathsf{FC}_*(\mathcal{S})$, contradicting the fact that the search $\mathcal{S}$ is successful. $\blacksquare$

Proposition 19. *For any connected graph G, $\mathsf{in}_m(G) \geqslant \mathsf{pw}(G) + 1$.*

Proof. Let $\mathcal{S}'$ be a successful monotonic search. We can modify each S_i' to S_i by removing the vertices which are not in $\mathsf{FC}_i(\mathcal{S}') \cup N_G(\mathsf{FC}_i(\mathcal{S}'))$, as well as removing the vertices in $\mathsf{FC}_{i-1}(\mathcal{S}')$. Note that $\mathsf{FC}_i(\mathcal{S}) = \mathsf{FC}_i(\mathcal{S}')$ for each i, so $\mathcal{S}$ is still successful and monotonic.

For each pair of adjacent vertices $\{v, w\}$, consider the minimal t for which after turn t, either v or w is in $\mathsf{FC}_t(\mathcal{S})$. Since $v, w \notin \mathsf{FC}_{t-1}(\mathcal{S})$ but $\{v, w\} \subseteq \mathsf{PC}_t(\mathcal{S})$, we must have $\{v, w\} \subseteq S_t$. Hence for each such pair there exists some S_t containing both. Now suppose $v \in S_i, S_j$ for some $i < j$. By the modification we performed, $v \notin \mathsf{FC}_{j-1}(\mathcal{S})$, and so by monotonicity $v \notin \mathsf{FC}_t(\mathcal{S})$ for any $t < j$. In particular, $v \notin \mathsf{FC}_i(\mathcal{S})$, yet $v \in S_i$. Hence, again by our modification, v has a neighbor in $\mathsf{FC}_i(\mathcal{S})$. This neighbor must remain fully cleared, so $v \in \mathsf{PC}_t(\mathcal{S})$ for each $t \geqslant i$. Hence $v \in S_t$ for each $i < t < j$, and so $\mathcal{S}$ is a path-decomposition of G of width $\max |S_i| - 1$, proving the claim. $\blacksquare$

These two propositions combined yield the following result:

Theorem 20. *For every connected graph G, $\mathsf{in}_m(G) = \mathsf{pw}(G) + 1$.*

Note that this implies that $\mathsf{in}(G) \leqslant \mathsf{pw}(G) + 1$. However, this inequality can be strict; in fact, we will see that the difference $\mathsf{in}_m(G) - \mathsf{in}(G)$ can be arbitrarily large.

5 The topological inspection number

Definition 21 (Subdivisions)**.** Given a graph G, an **edge subdivision** on an edge $\{v, w\} \in E(G)$ is the operation where we add a new vertex u, and replace the edge with a pair of

edges $\{v, u\}$ and $\{u, w\}$. A **subdivision** of G is a graph which can be obtained from G by a (possibly empty) finite sequence of edge subdivisions.

Example 22. There exist a graph G and a subdivision H of G such that $\mathsf{in}(H) \neq \mathsf{in}(G)$.

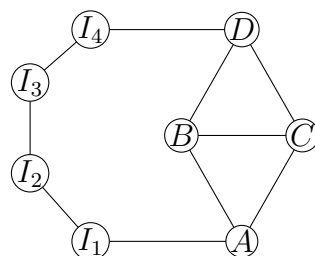*Proof.* Consider the subdivision of K_4 shown in Fig. 5.

Figure 5: A subdivision of K_4 with inspection number 3.

The following table lists a successful 3-search $\mathcal{S}$ of this graph:

i	S_i	$\mathsf{PC}_i(\mathcal{S})$	$\mathsf{FC}_i(\mathcal{S})$
1	$\{A, I_1, I_2\}$	$\{A, I_1, I_2\}$	$\{I_1\}$
2	$\{A, I_2, I_3\}$	$\{A, I_1, I_2, I_3\}$	$\{I_1, I_2\}$
3	$\{A, I_3, I_4\}$	$\{A, I_1, I_2, I_3, I_4\}$	$\{I_1, I_2, I_3\}$
4	$\{A, I_4, D\}$	$\{A, I_1, I_2, I_3, I_4, D\}$	$\{I_1, I_2, I_3, I_4\}$
5	$\{A, B, C\}$	$\{A, I_1, I_2, I_3, I_4, D, B, C\}$	$\{A, I_1, I_2, I_3\}$
6	$\{B, C, D\}$	$\{A, I_1, I_2, I_3, D, B, C\}$	$\{A, I_1, I_2, B, C\}$
7	$\{D, I_3, I_4\}$	$V(G)$	$V(G)$

However, $\mathsf{in}(K_4) = 4$ by Example 11, proving the claim. ∎

As explained in the introduction, we wish to investigate the "limiting" behavior of the inspection number under edge subdivisions, captured precisely in the following notion:

Definition 5 (Topological inspection number). The **topological inspection number** of a graph G, denoted $\mathsf{in}_t(G)$, is the minimum k such that for every subdivision H of G, there is a further subdivision H' of H with $\mathsf{in}(H') \leqslant k$.

Equivalently, $\mathsf{in}_t(G)$ is the minimum k such that for all ℓ, G has a subdivision H with $\mathsf{in}(H) \leqslant k$ in which every edge of G is subdivided at least ℓ times.

It is clear that $\mathsf{in}_t(H) = \mathsf{in}_t(G)$ for every subdivision H of G. We observe that $\mathsf{in}_t(G)$ is finite:

Proposition 23. $\mathsf{in}_t(G) \leqslant |V(G)| + 2$.

Proof. Let H be a subdivision of G. Enumerate the vertices of $V(G)$ as $v_1, \ldots, v_k$, and the edges in $E(G)$ as $e_1, \ldots, e_\ell$. Then each vertex in $V(H)$ is either in $V(G)$, or is the result of some series of subdivisions on an edge e_i. We may order these latter vertices in

order of the index of the subdivided edge they lie on, and break ties by distance to the lower-indexed endpoint of the subdivided edge (i.e., if $a, b \in V(H)$ lie on subdivided edge e_i with endpoints v_c, v_d with $c < d$, then whichever of a or b lies closer to v_c will come first in the order). Using this order we can enumerate the vertices in $V(H) \setminus V(G)$ as w_1, ..., w_m. Note that each w_i's neighbors are contained in $V(G) \cup \{w_{i-1}, w_{i+1}\}$.

Now consider the sequence of sets given by $S_t = V(G) \cup \{w_t, w_{t+1}\}$. Each w_i appears in either a single set or two consecutive sets while each v_i appears in every set. Furthermore, note that every pair of vertices with at least one vertex in $V(G)$ appears at least once, and every pair of vertices of the form $\{w_i, w_{i+1}\}$ appears once as well. Since every pair of neighbors must have one of these forms, every pair of neighbors appears together in a set of this sequence. This sequence thus witnesses $\mathsf{pw}(H) \leqslant |V(G)|+1$, and hence, by Theorem 20, $\mathsf{in}(H) \leqslant |V(G)| + 2$. Since H is a subdivision of itself, we have $\mathsf{in}_t(G) \leqslant |V(G)| + 2$. ∎

Evidently, the only connected graph G with $\mathsf{in}_t(G) = 1$ is the one-vertex graph K_1. Next we characterize the graphs G with $\mathsf{in}_t(G) = 2$:

Proposition 24. *If G is a connected graph with $\mathsf{in}_t(G) \leqslant 2$, then G is a path.*

Proof. Suppose G contains as a subgraph a cycle C. We wish to show that $\mathsf{in}(G) \geqslant 3$. Suppose otherwise and take a successful 2-search $\mathcal{S}$ of G. Let $v \in V(C)$ be such that the first index i with $v \in \mathsf{FC}_i(\mathcal{S})$ is minimal among all $v \in V(C)$. This implies that v and its two neighbors (which are distinct) in C are all contained in $\mathsf{PC}_i(\mathcal{S})$. However, by minimality of i, none of these vertices are contained in $\mathsf{FC}_{i-1}(\mathcal{S})$, and so all three must be contained in S_i, a contradiction. Since subdivisions of cyclic graphs remain cyclic, we must have G acyclic.

Now suppose, toward a contradiction, that G contains a vertex v_0 with $\deg(v_0) \geqslant 3$. Let G' be an arbitrary subdivision of G where each edge incident to v_0 is subdivided at least once. Let $v_1, v_2, v_3 \in N_{G'}(v_0)$ be distinct. Each of v_1, v_2, v_3 has a unique neighbor apart from v_0, and these neighbors are pairwise distinct. Let these neighbors be w_1, w_2, w_3 respectively. Let H be the subgraph of G' induced on $\{v_0, v_1, v_2, v_3, w_1, w_2, w_3\}$. It is straightforward to check that every subset $S \subseteq V(H)$ of size $|S| = 4$ has $|\partial_H(S)| \geqslant 2$. Hence, by Proposition 12 applied with $i = 5$, $\mathsf{in}(G') \geqslant \mathsf{in}(H) \geqslant 3$. Therefore, $\mathsf{in}_t(G) \geqslant 3$, which is a contradiction.

Hence, G must be an acyclic graph of maximum degree at most 2, i.e., a path. ∎

Since the inspection number of any path is at most 2 and subdivisions of paths are themselves paths, Proposition 24 implies that paths are precisely the connected graphs with topological inspection number at most 2. In the remainder of this paper we establish a characterization of the connected graphs G with $\mathsf{in}_t(G) = 3$, which turns out to be significantly more complicated.

6 Statement of the classification theorem

Definition 25 (Bipaths)**.** Let a **bipath** be the union of a pair of edge-disjoint paths P_1, P_2 which share at least three vertices, two of which are their endpoints, such that the shared

vertices appear on P_1 and P_2 in the same order. (See Fig. 6.)

We call the shared vertices the **primary vertices** of the bipath and the number of primary vertices the **order** of the bipath. The shared endpoints of P_1 and P_2 will be referred to as the **endpoints** of the bipath.

Note that each pair of consecutive primary vertices has a pair of internally vertex-disjoint paths between them contained in the bipath. These will be called **primary paths**.
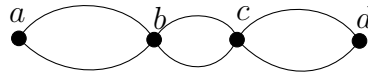

Figure 6: A bipath of order 4, with primary vertices a, b, c, d.

From here, we construct three families of graphs to be used in our theorem:

Definition 26 (Families $\boldsymbol{\mathcal{F}_1}$, $\boldsymbol{\mathcal{F}_2}$, $\boldsymbol{\mathcal{F}_3}$). Let $\mathcal{F}_1$ comprise all subdivisions of K_4.

Let $\mathcal{F}_2$ comprise the graphs consisting of 3 bipaths which share their endpoints but are otherwise pairwise vertex-disjoint.

Let $\mathcal{F}_3$ comprise the graphs consisting of 4 bipaths B_a, B_b, B_c, B_d and two paths P_1, P_2, such that B_a, B_b share endpoints v_1 and v_2, B_c, B_d share endpoints v_3 and v_4, but the bipaths are otherwise pairwise vertex-disjoint. Furthermore, P_1 has endpoints v_1, v_3, and P_2 has endpoints v_2, v_4, but these paths have no other vertices in common with the bipaths (however, P_1 and P_2 may intersect each other).

Examples of graphs from the families $\mathcal{F}_1$, $\mathcal{F}_2$, $\mathcal{F}_3$ are shown in Fig. 7.
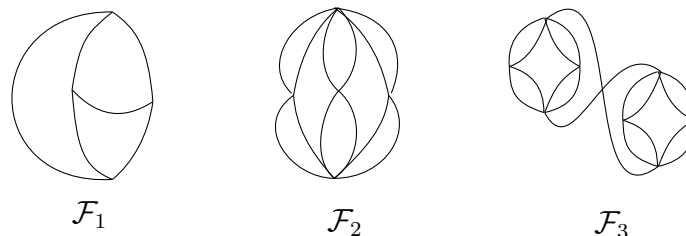

Figure 7: Representative members of the families $\mathcal{F}_1$, $\mathcal{F}_2$, $\mathcal{F}_3$. Each curved segment represents a path which may be subdivided arbitrarily many times.

Definition 27 (Generalized Series-Parallel Graphs). A **graph with terminals $\boldsymbol{G}$** is an ordered triple of the form (G, u, v) where G is a graph and u, v are distinct vertices in $V(G)$, called the **terminals** of $\boldsymbol{G}$. Note the convention of denoting graphs with terminals by the boldface version of the letter used to denote the underlying graph. The **generalized series-parallel graphs**, or **GSP graphs** for short, are the minimal collection of graphs with terminals that contains all single-edge graphs with their two vertices as terminals, and is closed under the following operations:

- The series operation $\circ_s$ can be performed on two GSP graphs $\boldsymbol{G} = (G, u, v)$ and $\boldsymbol{H} = (H, v, w)$ such that $V(G) \cap V(H) = \{v\}$. The result of this operation is the new GSP graph $\boldsymbol{G} \circ_s \boldsymbol{H} = (G \cup H, u, w)$.

- The parallel operation $\circ_p$ can be performed on two GSP graphs $\boldsymbol{G} = (G, u, v)$ and $\boldsymbol{H} = (H, u, v)$ such that $V(G) \cap V(H) = \{u, v\}$. The result of this operation is the new GSP graph $\boldsymbol{G} \circ_p \boldsymbol{H} = (G \cup H, u, v)$.

- The first branch operation $\circ_b$ can be performed on two GSP graphs $\boldsymbol{G} = (G, u, v)$, $\boldsymbol{H} = (H, u, w)$ such that $V(G) \cap V(H) = \{u\}$. The result of this operation is the new GSP graph $\boldsymbol{G} \circ_b \boldsymbol{H} = (G \cup H, u, v)$.

- The second branch operation $\circ_{b'}$ can be performed on two GSP graphs $\boldsymbol{G} = (G, u, v)$, $\boldsymbol{H} = (H, v, w)$ such that $V(G) \cap V(H) = \{v\}$. The result of this operation is the new GSP graph $\boldsymbol{G} \circ_{b'} \boldsymbol{H} = (G \cup H, u, v)$.

A graph G is **GSP** if there is some choice of terminals $u, v \in V(G)$ such that (G, u, v) is a GSP graph.

The way a GSP graph is built up from individual edges is encoded in its GSP decomposition:

Definition 28 (GSP Decompositions). Given a GSP graph $\boldsymbol{G} = (G, u, v)$, a **GSP decomposition** of $\boldsymbol{G}$ is a rooted binary tree T whose vertices are GSP graphs satisfying the following recursive conditions:

- The root node of T is $\boldsymbol{G}$.

- If G comprises a single edge, then T comprises a single node.

- Otherwise, there exist two GSP graphs $\boldsymbol{H}$ and $\boldsymbol{K}$ such that $\boldsymbol{G} = \boldsymbol{H} \circ \boldsymbol{K}$ for some operation $\circ \in \{\circ_s, \circ_p, \circ_b, \circ_{b'}\}$, and T comprises GSP decompositions T_H, T_K of $\boldsymbol{H}$ and $\boldsymbol{K}$, joined to the root node.

It is clear that every GSP graph $\boldsymbol{G}$ admits a GSP decomposition (not necessarily unique). Fig. 8 shows an example of a GSP decomposition.

Given a GSP decomposition T of $\boldsymbol{G}$ and $\boldsymbol{H} \in V(T)$, the **induced decomposition** of $\boldsymbol{H}$ is the subtree of T rooted at $\boldsymbol{H}$. If $\boldsymbol{H}$ is not a leaf of T, then $\circ_{T, \boldsymbol{H}}$ (or simply $\circ_{T, H}$, where H is the underlying graph of $\boldsymbol{H}$) denotes the operation in $\{\circ_s, \circ_p, \circ_b, \circ_{b'}\}$ by which $\boldsymbol{H}$ is obtained from its children.

We call T a **series-parallel decomposition** if $\circ_{T, \boldsymbol{H}} \in \{\circ_s, \circ_p\}$ for all $\boldsymbol{H} \in V(T)$.

Next we define the subclass of simple GSP decompositions.

Definition 29 (Simple GSP Decompositions). A graph with terminals is **bridged** if it contains an edge such that every path between its terminals includes that edge (i.e. a bridge separating its terminals).

We define the **complexity** $c(T)$ of a GSP decomposition T of $\boldsymbol{G}$ recursively as follows:
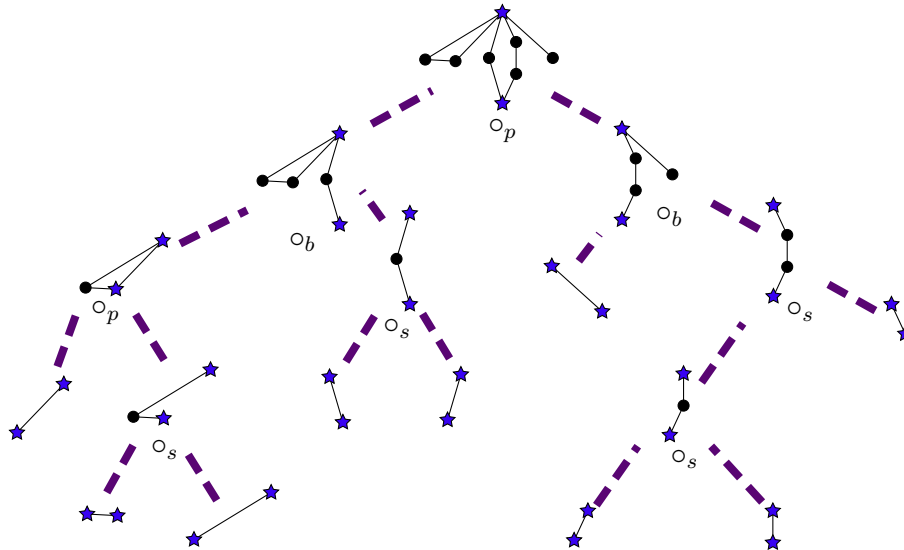
Figure 8: A GSP decomposition T. The starred vertices are terminals of their respective graphs. Each node $\boldsymbol{H}$ is labelled with $\circ_{T,\boldsymbol{H}}$.

- If $\circ_{T,\boldsymbol{G}} = \circ_p$, $c(T)$ is the sum of the complexities of the two subtrees rooted at the children of $\boldsymbol{G}$.

- If $\circ_{T,\boldsymbol{G}} \in \{\circ_b, \circ_{b'}\}$, $c(T)$ is the complexity of the child of $\boldsymbol{G}$ that shares its terminals with $\boldsymbol{G}$.

- Otherwise, the complexity of the decomposition is 0 if $\boldsymbol{G}$ is bridged and 1 otherwise.

A GSP decomposition T is **simple** if the complexity of every subtree of T is at most 1. Otherwise, it is **complex**.

Finally, we have all the terminology needed to state our result:

Theorem 30. *Let G be a connected graph with $|V(G)| \geqslant 2$. The following are equivalent:*

(1) $\mathsf{in}_t(G) \leqslant 3$.

(2) G *does not contain as a subgraph any element of* $\mathcal{F}_1, \mathcal{F}_2$, *or* $\mathcal{F}_3$.

(3) G *admits a simple GSP decomposition.*

The remainder of the paper is devoted to proving Theorem 30.

7 The topological inspection number of the forbidden graphs

7.1 A technical lemma

We begin by establishing a slightly technical result that will allow us to control sets of vertices with boundary of size at most 2 in our subsequent analysis.

Lemma 31. *Suppose G is a graph and $P_1, \ldots, P_k \subseteq V(G)$ are G-connected. For each i, let $Q_i = P_i \cup N_G(P_i)$, and let $P = \bigcup_i P_i$. Assume the following hypotheses:*

(H1) $\bigcup_{i=0}^{k} Q_i = V(G)$.

(H2) *For each $i \neq j$, $P_i \cap Q_j = \varnothing$.*

(H3) *For each i, $G - P_i$ is connected and $G - P_i - v$ is connected for every $v \in V(G) \setminus P$.*

(H4) *For each cut-set $\{a, b\} \subseteq P$ of G, there exists some i such that $\{a, b\} \subseteq P_i$.*

Then if $S \subseteq V(G)$ is such that $|\partial_G(S)| \leqslant 2$, at least one of the following conclusions holds:

(C1) $|S| \leqslant 2$;

(C2) $S \subseteq P_i \cup \{v\}$ *for some i and $v \in V(G) \setminus P$;*

(C3) $S \supseteq V(G) \setminus (P_i \cup A)$ *for some i and $A \subseteq V(G) \setminus P$ with $|A| \leqslant 1$;*

(C4) $S = \bigcup_{i \in I} P_i \cup A$ *for some $I \subseteq \{1, \ldots, k\}$ and $A \subseteq V(G) \setminus P$.*

Proof. Suppose $S \subseteq V(G)$ is a set with $|\partial_G(S)| \leqslant 2$. If $|S| \leqslant 2$, then (C1) holds. Thus, we may assume $|S| > 2$. Let $H = G - \partial_G(S)$. Then each connected component of H is either contained in S or is disjoint from S. Note that $V(H) \cap S \neq \varnothing$ since $|S| > |\partial_G(S)|$.

If H is connected then $S = V(G)$ and (C4) holds with $I = \{1, \ldots, k\}$ and $A = V(G) \setminus P$. Thus, we may assume H is disconnected.

Suppose $\partial_G(S) \subseteq V(G) \setminus P$. For every i, $\partial_G(S)$ is disjoint from P_i and hence P_i is either contained in S or disjoint from S. Letting $I = \{i \mid P_i \subseteq S\}$ and $A = S \setminus P$ gives us $S = \bigcup_{i \in I} P_i \cup A$, so (C4) holds.

Otherwise, $|\partial_G(S) \setminus P| \leqslant 1$. Note that we can find some i such that $\partial_G(S) \cap P \subseteq P_i$, as $\partial_G(S) \cap P$ either comprises a single vertex or $\partial_G(S) \subseteq P$ and we can use (H4). By (H3), the graph $G - P_i - \partial_G(S) = G - P_i - (\partial_G(S) \setminus P)$ is connected and is thus either contained in S or disjoint from S. In the former case (C3) holds, and in the latter case S is contained in $P_i \cup (\partial_G(S) \setminus P)$, fulling (C2). $\blacksquare$

7.2 The family $\mathcal{F}_1$

Here we show that the topological inspection number of K_4 is 4:

Theorem 32. $\mathsf{in}_t(K_4) \geqslant 4$.

The remainder of this subsection contains the proof of Theorem 32. Note that Theorem 32 implies that every graph in $\mathcal{F}_1$ (i.e., every subdivision of K_4) has topological inspection number at least 4.

Let G be some subdivison of K_4 such that each edge is subdivided at least 2^{12} times (we did not make an attempt to optimize the value 2^{12}). We will show that $\mathsf{in}(G) \geqslant 4$. Let

$T = \{a, b, c, d\}$ be the original vertices of the K_4 in this subdivision. Assume, towards a contradiction, that $\mathsf{in}(G) \leqslant 3$.

We begin by applying Lemma 31 in the context of G. Let P_{ab}, P_{ac}, ..., P_{cd} be the sets of internal vertices of the paths replacing the corresponding edges of K_4. In the notation of Lemma 31, each set Q_{xy} would then be of the form $Q_{xy} = P_{xy} \cup \{x, y\}$.

Claim 33. *The sets P_{ab}, P_{ac}, ..., P_{cd} satisfy the assumptions of Lemma 31.*

Proof. It is easy to see that the sets P_{xy} are G-connected and that statements (H1) and (H2) hold. For each pair (x, y), $G - P_{xy}$ is a subdivision of a diamond graph, which is 2-connected, so (H3) is satisfied. It thus remains to verify (H4).

If $\{v, w\}$ is a cut-set such that $\{v, w\} \subseteq P$—i.e., $\{v, w\} \cap T = \varnothing$—then there exists a path through all 4 vertices of T which passes through neither v nor w. Any vertex not on this path lies on a subdivided edge between two vertices of T. For such a vertex to be separated from T by $\{v, w\}$, both v and w must be on this subdivided edge, and thus they both belong to some P_{xy}, as desired. ∎

Now consider the possible sizes of subsets of $V(G)$ with boundary of size at most 2. Let $M = \max |P_{xy}|$. Suppose $S \subseteq V(G)$ is a set with $|\partial(S)| \leqslant 2$ and $M + 1 < |S| < |V(G)| - M - 1$. Then, by Lemma 31,

$$S = \bigcup_{xy \in I} P_{xy} \cup T' \quad \text{for some } I \subseteq \{ab, ac, \ldots, cd\} \text{ and } T' \subseteq T.$$

There are at most 2^{10} such sets S (as we have 2^4 options for T' and 2^6 options for I). On the other hand, by Proposition 12, for every $i \in (M + 1, |V(G)| - M - 2)$, we must have such a set S with $|S| = i$ or $|S| = i + 1$. Hence, $|V(G)| - 2M - 4 \leqslant 2 \cdot 2^{10}$, which yields

$$|V(G)| \leqslant 2M + 2^{11} + 4. \tag{§}$$

Now suppose without loss of generality that $|P_{ab}| = M$. Let

$$R_1 = \{a\} \cup P_{ac} \cup P_{ad}, \quad R_2 = \{b\} \cup P_{bc} \cup P_{bd}, \quad \text{and} \quad R_3 = P_{cd}.$$

Without loss of generality, suppose $|R_1| \leqslant |R_2|$. Let $N = \max_{i=1,2,3} |R_i|$. Note that $N \leqslant |V(G)| - M$. Since every edge of K_4 is subdivided at least 2^{12} times, we have $|R_1| \geqslant 2^{12}$, so, by (§),

$$N - |R_1| \leqslant |V(G)| - M - 2^{12} \leqslant M - 2^{11} + 4 < M.$$

Let D be the induced subgraph of G on the union of $R_1 \cup R_2 \cup R_3 \cup \{c, d\}$. Let R_4 be the set of the $N - |R_1|$ vertices in P_{ab} that are closest to a, and let $H = G[V(D) \cup R_4]$ (see Fig. 9).

Claim 34. *The sets R_1, R_2, R_3 satisfy the assumptions of Lemma 31 as subsets of D.*
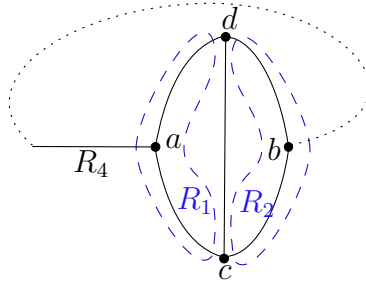
Figure 9: Construction of the subgraph H.

Proof. It is easy to see that the set R_i are connected and that statements (H1) and (H2) hold. For each i, $D - R_i$ is a cycle, which is 2-connected, so (H3) holds. It thus remains to show (H4).

Note that no set $\{x, y\} \subseteq R_1 \cup R_2 \cup R_3$ may separate c from d, as there are 3 internally disjoint paths between them. Hence for $\{x, y\}$ to be a cut-set of D, it must separate some vertex v from both c and d. Whichever R_i contains v gives two internally disjoint paths from v to each of c, d, so one of x, y must lie on each of these paths. Hence, x and y are both contained in R_i, as desired. ∎

Claim 35. *Suppose a set $S \subseteq V(H)$ satisfies $|\partial_H(S)| \leqslant 2$. Then $|S| \leqslant N + 2$ or $|S| \geqslant |V(H)| - N - 2$.*

Proof. Let $S' = S \cap V(D)$. Note that $|\partial_D(S')| \leqslant 2$ as any vertex in $\partial_D(S')$ is in $\partial_H(S)$ as well.

Since D is 2-connected, either $|S'| < 2$, or $V(D) \subseteq S$, or $|\partial_D(S')| \geqslant 2$ (by Lemma 8). In the first case, $|S| \leqslant N - |R_1| + 1 < N + 2$. In the second case,

$$|S| \geqslant |V(D)| = |V(H)| - N + |R_1| > |V(H)| - N - 2.$$

The remaining case is where $|\partial_D(S')| = 2$. This implies that $\partial_H(S) = \partial_D(S')$. In particular, R_4 does not contain a boundary vertex of S. Since R_4 is connected, it is either entirely contained in S or is disjoint from it. Note that if $R_4 \subseteq S$, then $a \in S$ as otherwise the neighbor of a in R_4 will be a boundary vertex, a contradiction. Also, if $S \cap R_4 = \varnothing$, then either $a \notin S$ or $a \in \partial_D(S')$.

By Lemma 31, one of the following four cases must hold.

Case 1: $|S'| \leqslant 2$. Then $|S| \leqslant N - |R_1| + 2 < N + 2$.

Case 2: S' *is a subset of* $R_i \cup \{c, d\}$ *for some* $1 \leqslant i \leqslant 3$. If $i = 1$, this implies $|S| \leqslant N - |R_1| + |R_1| + 2 = N + 2$. Otherwise, $a \notin S$ so S is disjoint from R_4. Then $|S| \leqslant \max\{|R_2|, |R_3|\} + 2 \leqslant N + 2$.

Case 3: S' *contains* $V(D) \setminus (R_i \cup \{c, d\})$ *for some* $1 \leqslant i \leqslant 3$. If $i \neq 1$, then $R_1 \subseteq S$, which implies that a is in S but not in $\partial_D(S')$. Therefore, $R_4 \subseteq S$, so $|S| \geqslant |V(H)| - \max\{|R_2|, |R_3|\} - 2 \geqslant |V(H)| - N - 2$. Otherwise, $|S| \geqslant |V(H)| - (|R_1| + |R_4| + 2) = |V(H)| - N - 2$.

Case 4: *S' is the union of some subset of $\{R_1, R_2, R_3, \{c\}, \{d\}\}$.* If $R_2 \cup R_3 \subseteq S$, then $|S| \geqslant |V(H)| - N - 2$. On the other hand, if both $R_2, R_3 \not\subseteq S$, then $|S| \leqslant |R_1| + |R_4| + 2 = N + 2$. The only remaining possibility is that exactly one of R_2, R_3 is a subset of S. Note that the cycle formed by $R_2 \cup R_3 \cup \{c, d\}$ contains two boundary vertices of S. Thus $R_1 \subseteq S$ if and only if $R_4 \subseteq S$, as otherwise one of a or its neighbor in R_4 would be a third boundary vertex. If $R_1 \not\subseteq S$, then $|S| \leqslant \max\{|R_2|, |R_3|\} + 2 = N + 2$. If $R_1 \subseteq S$, then $|S| \geqslant |R_1| + \min\{|R_2|, |R_3|\} + |R_4| \geqslant |V(H)| - N - 2$. ∎

Note that
$$|V(H)| \geqslant 2N + \min\{|R_2|, |R_3|\} + 2 \geqslant 2N + 2^{12} + 2,$$
so $|V(H)| - N - 2 \geqslant N + 2^{12}$. Hence there does not exist a set $S \subseteq V(H)$ with $|\partial_H(S)| \leqslant 2$ and $N + 2 < |S| < N + 2^{12}$, and so by Proposition 12, $\mathsf{in}(G) \geqslant \mathsf{in}(H) \geqslant 4$, a contradiction.

Therefore, each subdivision of K_4 where every edge is subdivided at least 2^{12} times has inspection number at least 4, which implies $\mathsf{in}_t(K_4) \geqslant 4$.

7.3 The families $\mathcal{F}_2$ and $\mathcal{F}_3$

To show the bound $\mathsf{in}_t \geqslant 4$ for graphs in $\mathcal{F}_2$ and $\mathcal{F}_3$, we first consider a configuration common to both families. Suppose G is a graph comprising two vertices, a, b, two internally disjoint bipaths between them (the union of which we call D), and a path with a as an endpoint which is otherwise disjoint from D (see Fig. 10). We let R be the set of the vertices on this path distinct from a. Let I be the set of all pairs $\{x, y\}$ of consecutive primary vertices on the bipaths forming D. For each $\{x, y\} \in I$, let $P_{xy,0}$ and $P_{xy,1}$ be the sets of internal vertices on the two primary paths joining x and y. We assume that all the sets $P_{xy,i}$ are nonempty.
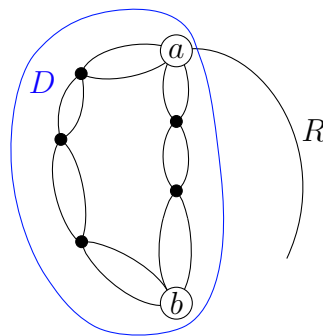
Figure 10: The special configuration H.

Claim 36. *The sets $P_{xy,i}$ satisfy the hypotheses of Lemma 31 applied to the graph D.*

Proof. Evidently, the sets $P_{xy,i}$ are connected and both (H1), (H2) hold. Removing any set $P_{xy,i}$ from D keeps the remaining graph 2-connected, showing (H3). Finally, if v and w are two non-primary vertices that do not lie on the same primary path, then $D - \{v, w\}$ is connected, proving (H4). ∎

Now we are going to make a few additional assumptions about the graph H. First, let T be the sum of the orders of the bipaths comprising D (recall that the order of a bipath is the number of its primary vertices). We shall assume that:

- $|P_{xy,i}| \geqslant 2^{4T}$ for all $\{x, y\} \in I$ and $i \in \{0, 1\}$.

Next, let N be the maximum of $|P_{xy,i}|$ taken over all sets $P_{xy,i}$ with $a \notin \{x, y\}$. Similarly, let M be the maximum of $|P_{xy,i}|$ taken over all sets $P_{xy,i}$ with $a \in \{x, y\}$. We assume that:

- $N \geqslant M$ and $|R| = N - M$.

With these assumptions, we can show that the inspection number of H is at least 4:

Proposition 37. $\mathsf{in}(H) \geqslant 4$.

Proof. Suppose $S \subseteq V(H)$ is a subset such that $|\partial_G(H)| \leqslant 2$ and

$$N + 2 \leqslant |S| \leqslant |V(H)| - N - 2.$$

Let $S' = S \cap V(D)$. Note that $\partial_D(S') \subseteq \partial_H(S)$. Since D is 2-connected, we have $|S'| < 2$, or $V(D) \subseteq S$, or $|\partial_D(S')| = 2$ (by Lemma 8). In the first case, $|S| \leqslant |R|+1 = N-M+1 < N + 2$, a contradiction. In the second case,

$$|S| \geqslant |V(D)| = |V(H)| - |R| = |V(H)| - N + M > |V(H)| - N - 2,$$

which is again a contradiction. Therefore, $|\partial_D(S')| = 2$, which implies that $\partial_H(S) = \partial_D(S')$. In particular, R does not contain any boundary vertices of S. Since R is H-connected, it is either entirely contained in S or is disjoint from it. Note that if $R \subseteq S$, then $a \in S$ as otherwise the neighbor of a in R will be a boundary vertex, a contradiction. Also, if $S \cap R = \varnothing$, then either $a \notin S$ or $a \in \partial_D(S')$.

Since $|\partial_D(S')| \leqslant 2$, by Lemma 31, one of the following four cases must hold.

Case 1: $|S'| \leqslant 2$. Then $|S| \leqslant |R| + 2 = N - M + 2 < N + 2$, a contradiction.

Case 2: S' is contained in $P_{xy,i} \cup \{z\}$ for some $P_{xy,i}$ and a primary vertex z. If $a \in \{x, y\}$, then $|P_{xy,i}| \leqslant M$, so $|S| \leqslant |R| + M + 1 = N + 1$, a contradiction. Thus, $a \notin \{x, y\}$. Furthermore, $z = a$, as otherwise $R \cap S = \varnothing$, which again implies $|S| \leqslant N + 1$. Therefore, $a \in \partial_H(S)$. Consider the cycle C in H with $V(C) = P_{xy,0} \cup P_{xy,1} \cup \{x, y\}$. It is 2-connected, not contained in S, and, since $a \notin V(C)$, it contains at most one boundary vertex of S. Thus, $|S \cap V(C)| \leqslant 1$, so $|S| \leqslant |R| + 2 < N + 2$, a contradiction.

Case 3: S' contains $V(D) \setminus (P_{xy,i} \cup \{z\})$ for some $P_{xy,i}$ and a primary vertex z. Then S contains at least 3 sets $P_{uv,j}$ with $a \in \{u, v\}$. If $a \notin S$, each of these sets includes a boundary vertex of S, a contradiction. Hence $a \in S$. If $a \in \{x, y\}$, then $|S| \geqslant |V(H)| - M - |R| - 1 = |V(H)| - N - 1$. Otherwise, considering the cycle C in H with vertex set $P_{xy,0} \cup P_{xy,1} \cup \{x, y\}$ as we did in Case 2 yields that $a \notin \partial_H(S)$. Hence $R \subseteq S$, so $|S| \geqslant |V(H)| - N - 1$, a contradiction.

Therefore, any such set S has to satisfy the following:

Case 4: *S′ is the union of a collection of sets $P_{xy,i}$ and some primary vertices.* Since there are $2T-4$ sets $P_{xy,i}$ and $T-2$ primary vertices, the number of such sets S is at most

$$2^{2T-4} \cdot 2^{T-2} \cdot 2 = 2^{3T-5},$$

where we are multiplying by 2 to account for the fact that S either contains R or is disjoint from R. If $\mathsf{in}(H) \leqslant 3$, then, by Proposition 12, for each $i \in [N+2, |V(H)| - N - 2)$, there exists some $S \subseteq V(H)$ with $|\partial_H(S)| \leqslant 2$ and $|S| \in \{i, i+1\}$. Therefore,

$$|V(H)| - 2N - 4 \leqslant 2 \cdot 2^{3T-5} = 2^{3T-4} < 2^{3T}.$$

However, $V(H)$ contains as disjoint subsets R, a set $P_{xy,i}$ with $a \in \{x, y\}$ of size M, a set $P_{xy,i}$ with $a \notin \{x, y\}$ of size N, and at least one other set $P_{xy,i}$, so

$$|V(H)| \geqslant (N-M) + M + N + 2^{4T} > 2N + 4 + 2^{3T},$$

a contradiction. ∎

Theorem 38. *Suppose $G_0 \in \mathcal{F}_2$. Then $\mathsf{in}_t(G_0) \geqslant 4$.*

Proof. Let a, b be the shared vertices of the three bipaths forming G_0 (see Definition 26). Let T be the sum of the orders of the bipaths comprising G_0, and suppose G is some subdivision of G_0 such that each edge is subdivided at least 2^{4T} times. We will show that $\mathsf{in}(G) \geqslant 4$. Note that G is in $\mathcal{F}_2$ and that every primary path in G contains at least 2^{4T} internal vertices. Let P be a primary path of maximum length in G. Let D be the union of the two bipaths not containing P. Let M (resp. L) be the maximum number of internal vertices on a primary path in D adjoining a (resp. b). Without loss of generality, assume that $L \geqslant M$ and let N be the maximum number of internal vertices on a primary path in D not adjoining a. Note that $N \geqslant L \geqslant M$ and P has at least N internal vertices.

Let Q be an arbitrary ab-path through P disjoint from D. Let R be the set of $N-M$ vertices in $V(Q) \setminus \{a\}$ closest to a. The induced subgraph H of G on $V(D) \cup R$ satisfies the assumptions of Proposition 37. Therefore, $\mathsf{in}(G) \geqslant \mathsf{in}(H) \geqslant 4$, as desired. ∎

Theorem 39. *Suppose $G_0 \in \mathcal{F}_3$. Then $\mathsf{in}_t(G_0) \geqslant 4$.*

Proof. Recall that by Definition 26, G_0 consists of 4 bipaths B_a, B_b, B_c, B_d and two paths P_1, P_2, such that B_a, B_b share endpoints v_1 and v_2, B_c, B_d share endpoints v_3 and v_4, but the bipaths are otherwise pairwise vertex-disjoint. Furthermore, P_1 has endpoints v_1, v_3, and P_2 has endpoints v_2, v_4, but these paths have no other vertices in common with the bipaths (however, P_1 and P_2 may intersect each other).

Let T be the sum of the orders of the bipaths B_a, B_b, B_c, and B_d, and suppose G is some subdivision of G_0 such that each edge is subdivided at least 2^{4T} times. We will show that $\mathsf{in}(G) \geqslant 4$. Note that G is in $\mathcal{F}_3$ and every primary path in G has at least 2^{4T} internal vertices. Let P be a primary path of maximum length in G. Suppose without loss of generality that P belongs to a bipath with endpoints v_3, v_4. Let D be the subgraph of G comprising the two bipaths with endpoints v_1, v_2. Let M (resp. L) be the maximum

number of internal vertices on a primary path in D adjoining v_1 (resp. v_2). Without loss of generality, assume that $L \geqslant M$ and let N be the maximum number of internal vertices on a primary path in D not adjoining v_1. Note that $N \geqslant L \geqslant M$ and P has at least N internal vertices.

Let Q be a path that starts at v_1, follows the path P_1 to v_3, and then continues to v_4 through P. Let R be the set of $N - M$ vertices in $V(Q) \setminus \{v_1\}$ closest to v_1. The induced subgraph H of G on $V(D) \cup R$ satisfies the assumptions of Proposition 37 (with v_1 in place of a and v_2 in place of b). Therefore, $\mathsf{in}(G) \geqslant \mathsf{in}(H) \geqslant 4$, as desired. ∎

8 Building simple GSP decompositions

8.1 GSP decompositions of K_4-subdivision-free graphs

The main result of this subsection is Corollary 44, which says that connected K_4-subdivision-free graphs admit GSP decompositions (with some control over the choice of terminals). By Theorem 32, every graph G with $\mathsf{in}_t(G) \leqslant 3$ must be K_4-subdivision-free, so Corollary 44 implies that such G has a GSP decomposition. In the subsequent subsections we shall explain how to make this GSP decomposition simple.

Lemma 40. *Suppose G is a K_4-subdivision-free graph and H is a connected subgraph of G containing distinct vertices a, b such that $a \notin N_H(b)$, $\{a, b\}$ is not a cut-set of H, and there exists a path from a to b in $G - (H - \{a, b\})$. Then there exists a cut vertex of H separating a and b.*

Proof. Suppose H contained internally disjoint paths P, Q between a and b. Since $a \notin N_H(b)$, P and Q must each contain at least one intermediate vertex. As H remains connected after removing a and b, there must be some path connecting intermediate vertices of P and Q that is internally disjoint from P and Q. Let v and w be the endpoints of such a path. Then there are pairwise disjoint paths connecting each of the pairs of vertices in $\{a, b, v, w\}$ (the path between a and b in particular is taken to be one contained in $G - (H - \{a, b\})$ given by the hypothesis). Thus G contains a subdivision of K_4, a contradiction. Hence H cannot contain a pair of internally disjoint paths between a and b, and so the claim follows by Menger's theorem [10, Corollary 3.3.5]. ∎

Proposition 41. *Suppose G is a 2-connected K_4-subdivision-free graph, and $\{a, b\} \subseteq V(G)$ are either neighbors or a cut-set of G. Then (G, a, b) admits a series-parallel decomposition.*

Proof. We shall prove this via induction on $|V(G)|$. The base case where $|V(G)| = 2$ is trivial.

To prove the general case, let the connected components of $G - \{a, b\}$ be $H'_0, \ldots, H'_k$, and for each i, let H_i be obtained from $G[H'_i \cup \{a, b\}]$ by removing the edge between a and b if it exists. For each i, the graph $G - H'_i$ contains an ab-path P_i. Indeed, if a and b are neighbors, the edge between a and b is such a path. Otherwise, $\{a, b\}$ is a cut-set of G so there are at least two connected components of $G - \{a, b\}$. Hence $k > 0$ and we can

choose some H'_j distinct from H'_i. Since G is 2-connected, $G[H'_j \cup \{a, b\}]$ is a connected subgraph of $G - H'_i$, and thus it contains an ab-path. By Lemma 40, we may conclude that H_i must contain a cut-vertex separating a and b.

Consider the block-cut tree of H_i. If it is not a single path from the block containing a to the block containing b (which are distinct by the above), we can find some leaf block L with a cut-vertex v such that $a, b \notin V(L) \setminus \{v\}$. Then v separates $V(L) \setminus \{v\}$ from the rest of G, which is impossible as G is 2-connected.

Thus the block-cut tree of H_i is a path from the block containing a to the block containing b. Enumerate the blocks of H_i as $B_0, \dots, B_{\ell-1}$ in order of increasing distance from a, and let v_j be the cut-vertex between B_{j-1} and B_j. We also write $v_0 = a$, $v_\ell = b$. We can find a path from each v_j to v_{j+1} without passing through any other vertices of B_j: take a path from v_j to a through the blocks $B_{j-1}, \dots, B_0$ in order, then P_i, then a path from b to v_{j+1} through the blocks $B_{\ell-1}, \dots, B_{j+1}$ in order. Furthermore, B_j cannot contain a cut-vertex separating v_j from v_{j+1}, so by the contrapositive of Lemma 40, either $\{v_j, v_{j+1}\}$ is a cut-set of B_j, or $v_j \in N_{B_j}(v_{j+1})$. In either case we can apply the induction hypothesis to see that (B_j, v_j, v_{j+1}) admits a series-parallel decomposition.

Note that

$$(H_i, a, b) = (B_0, v_0, v_1) \circ_s (B_1, v_1, v_2) \circ_s \cdots \circ_s (B_{\ell-1}, v_{\ell-1}, v_\ell),$$

so this implies that (H_i, a, b) admits a series-parallel decomposition. Furthermore, we have

$$(G, a, b) = \begin{cases} (H_0, a, b) \circ_p \cdots \circ_p (H_k, a, b) & \text{if } a \notin N_G(b); \\ (H_0, a, b) \circ_p \cdots \circ_p (H_k, a, b) \circ_p (e_{ab}, a, b) & \text{if } a \in N_G(b), \end{cases}$$

where e_{ab} is the graph comprising the single edge between a and b, so G admits a series-parallel decomposition, as desired. ∎

Lemma 42. *If $\boldsymbol{G} = (G, a, b)$ admits a series-parallel decomposition T, then G's block-cut tree comprises a single path (possibly a single block), and a and b are each contained in a leaf block. If we let $v_1, \dots, v_{k-1}$ be the cut-vertices of G in order of distance from a, and $v_0 = a$, $v_k = b$, then for every block B, $(B, v_i, v_{i+1}) \in V(T)$ where i is the unique index such that $\{v_i, v_{i+1}\} \subseteq V(B)$. Furthermore, for any $v \in V(G)$, we can find a path from a through v to b.*

Proof. We shall show this by induction on $|V(T)|$. The claim is clearly true if T is trivial, where G must comprise a single edge.

If $\circ_{T, \boldsymbol{G}} = \circ_s$, suppose (H, a, c) and (K, c, b) are the children of $\boldsymbol{G}$ in T. Then c separates $V(H) \setminus \{c\}$ from $V(K) \setminus \{c\}$ in G, so the block-cut tree of G is simply the union of the block-cut trees of H and K with an additional edge between each of their blocks which contain c. One notes that this, combined with the induction hypothesis, gives the desired structure of the block-cut graph of G. Each block (B, v_i, v_{i+1}) is either a block of H or a block of K, so it is in $V(T)$ by the induction hypothesis. Finally, for any $v \in V(G)$, suppose without loss of generality that $v \in V(H)$. Then taking a path in H from a through

v to c and extending it by any path in K from c to b gives a desired path from a through v to b in G.

If $\circ_{T,\boldsymbol{G}} = \circ_p$, we will show that G is 2-connected. Let (H, a, b) and (K, a, b) be the children of $\boldsymbol{G}$ in T. Consider any vertex $v \in V(G)$. We will show that the graph $G - v$ is connected. We may assume that $v \neq a$ (the argument when $v = a$ is the same but with the roles of a and b exchanged). We claim that every vertex $w \in V(G) \setminus \{v\}$ is reachable in $G - v$ by a path from a. Without loss of generality, we may assume that $w \in V(H)$. By the induction hypothesis, we can find a path P from a through w to b in H and a path Q from a to b in K. Write $P = P_1 \cup P_2$, where P_1 is an aw-path and P_2 is a wb-path, noting that P_1 and P_2 share only a single vertex w. If $v \notin V(P_1)$, then P_1 is a desired aw-path in $G - v$. Otherwise, i.e., if $v \in V(P_1)$, a desired aw-path in $G - v$ is $Q \cup P_2$. Now that we have established that G is 2-connected, the block-cut graph of G comprises a single block, and the desired properties in the claim follow immediately. To find a path from a through v to b for any $v \in V(G)$, we simply take the path in whichever of H or K contains v using the induction hypothesis. ∎

Lemma 43. *Suppose $\boldsymbol{G} = (G, a, b)$ and $\boldsymbol{H} = (H, c, d)$ admit GSP decompositions T_G, T_H respectively, and $V(G) \cap V(H) = \{c\}$. Let K be the union of G and H. Then $\boldsymbol{K} = (K, a, b)$ admits a GSP decomposition T with $\boldsymbol{H} \in V(T)$ such that the induced decomposition of $\boldsymbol{H}$ in T is T_H.*

Furthermore, if both T_G and T_H are simple, then we can additionally further restrict T to be simple and have the same complexity as T_G.

Proof. We shall proceed by induction on $|V(G)|$. For the base case, $|V(G)| = 2$ and G comprises a single edge. If $c = a$, $\boldsymbol{K} = \boldsymbol{G} \circ_b \boldsymbol{H}$. Similarly, if $c = b$, $\boldsymbol{K} = \boldsymbol{G} \circ_{b'} \boldsymbol{H}$ and T can be obtained directly by joining T_G to T_H. In either case, the complexity of T is equal to the complexity of T_G, so if T_G and T_H were simple, T would be simple too.

In the general case, T_G is nontrivial. Let $\boldsymbol{G_0} = (G_0, s, t)$, $\boldsymbol{G_1}$ be the children of $\boldsymbol{G}$ in T_G. Without loss of generality we may assume that $c \in V(G_0)$. Then by the induction hypothesis, $(G_0 \cup H, s, t)$ admits a GSP decomposition T' containing $\boldsymbol{H}$ with induced decomposition T_H. Connecting T' with the induced decomposition of T_G on $\boldsymbol{G_1}$ gives the desired decomposition T. If T_G and T_H were simple, the induced decompositions on G_0 and G_1 would be simple too. By the induction hypothesis, T' is simple and has the same complexity as the induced decomposition on G_0 by T_G, so T has the same complexity as T_G and is simple. ∎

Corollary 44. *Suppose G is a connected K_4-subdivision-free graph, and $\{a, b\} \subseteq V(G)$ are either neighbors or a cut-set of a 2-connected component of G. Then (G, a, b) is a GSP graph.*

Proof. We proceed by induction on the number of blocks in G. In the case where the number of blocks is 1, this reduces to Proposition 41. In the general case, take a leaf block L which does not contain both a and b. Let v be the cut-vertex of L. By the induction hypothesis, $(G - (L - \{v\}), a, b)$ is a GSP graph, and by Proposition 41, $(L, \{v, w\})$ is a GSP graph for any neighbor w of v in L. Hence by Lemma 43, (G, a, b) is a GSP graph. ∎

8.2 Simple series-parallel decompositions of 2-connected graphs

The bulk of the work in the proof of the implication (2) $\implies$ (3) in Theorem 7 is done in this subsection, where we show that 2-connected graphs G without any member of $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$ as a subgraph admit simple series-parallel decompositions (Proposition 52).

Lemma 45. *Suppose $\boldsymbol{G} = (G, a, b)$ has a GSP decomposition T with complexity k. Then G contains at least k pairwise internally vertex-disjoint bipaths with endpoints $\{a, b\}$.*

Proof. We proceed by induction on $|V(G)|$. If T is trivial, we are done as then $k = 0$.

If $\circ_{T, \boldsymbol{G}} = \circ_p$, $\circ_b$ or $\circ_{b'}$, then by definition k is the sum of the complexities of the children of $\boldsymbol{G}$ which have terminals (a, b). By the induction hypothesis, each of these children have at least as many bipaths with terminals $\{a, b\}$ as their respective complexities, and since these children may only intersect at $\{a, b\}$, these bipaths are internally vertex-disjoint, proving the claim.

Otherwise, $\circ_{T, \boldsymbol{G}} = \circ_s$. Then, by definition, either $\boldsymbol{G}$ is bridged and $k = 0$ or $\boldsymbol{G}$ is not bridged and $k = 1$. In the former case, there is nothing to prove, so we may suppose that $k = 1$. Consider the block-cut tree of G. Since $\circ_{T, \boldsymbol{G}} = \circ_s$, there exists at least one cut-vertex separating a from b, so a and b do not share a block. Let P be the unique path in the block-cut tree from the block containing a to the block containing b. If any block in P comprises a single edge, this edge would be a bridge separating a from b, contradicting the assumption that $\boldsymbol{G}$ is not bridged. Hence every block in P contains at least 3 vertices, and thus has a cycle containing the cut-vertices it shares with its neighbors in P (or in the case of the endpoint blocks, the one shared cut-vertex and whichever of a or b it contains). The union of these cycles is a bipath, as desired. ∎

We can note the following corollary, which immediately results from the previous lemma:

Corollary 46. *Suppose T is a complex GSP decomposition of a graph G. Then some $\boldsymbol{H} \in V(T)$ contains as a subgraph a pair of internally disjoint bipaths between its terminals.*

Lemma 47. *Suppose $\boldsymbol{G} = (G, a, b)$ admits a simple GSP decomposition T. Then (G, b, a) also admits a simple GSP decomposition. Furthermore, if T is a series-parallel decomposition then (G, b, a) also admits a simple series-parallel decomposition.*

Proof. We proceed by induction on $|V(G)|$. The base case where G comprises a single edge is trivial. Suppose (G_0, s_0, t_0) and (G_1, s_1, t_1) are the children of $\boldsymbol{G}$ in T. Their induced decompositions must be simple GSP decompositions, so by the induction hypothesis (G_0, t_0, s_0), (G_1, t_1, s_1) also admit simple GSP decompositions. If $\circ_{T, \boldsymbol{G}} = \circ_s$ or $\circ_p$, then $\boldsymbol{G} = (G_1, t_1, s_1) \circ_{T, \boldsymbol{G}} (G_0, t_0, s_0)$. Otherwise, $\boldsymbol{G} = (G_1, t_0, s_0) \circ_* (G_0, s_1, t_1)$ where $\circ_*$ is whichever of $\{\circ_b, \circ_{b'}$ is not $\circ_{T, \boldsymbol{G}}$. In either case, connecting the two GSP decompositions given of the children gives a GSP decomposition of G as desired. ∎

Proposition 48. *Suppose G is a 2-connected graph and $(G, a, b) = (H, a, b) \circ_p (K, a, b)$, where both (H, a, b) and (K, a, b) admit simple series-parallel decompositions T_H, T_K respectively. Then G admits a simple series-parallel decomposition with a as a terminal.*

Proof. Suppose without loss of generality that $|V(H)| \leqslant |V(K)|$. We argue by induction on $|V(H)| + |E(H)|$.

For the base case, where $|V(H)| + |E(H)| = 3$, i.e. $H \cong K_2$, directly combining T_H and T_K gives a series-parallel decomposition T of G with complexity at most 1, since H comprises a single edge and so T_H is trivial, with complexity 0. Every proper subtree of T is a subtree of T_H or T_K, and thus must have complexity at most 1, so T is simple as desired.

In the general case, T_H is nontrivial. Suppose $\circ_{T_H, \boldsymbol{H}} = \circ_p$ and (H_0, a, b), (H_1, a, b) are the children of $\boldsymbol{H}$ in T. Since T_H has complexity at most 1, one of these children has an induced decomposition of complexity 0; suppose without loss of generality that this is H_1. This naturally gives a simple series-parallel decomposition of $((H_1, a, b) \circ_p (K, a, b))$. Then, note that

$$(G, a, b) = (H_0, a, b) \circ_p ((H_1, a, b) \circ_p (K, a, b)).$$

Since H_0 is a proper subgraph of H, the induction hypothesis implies that G has a simple series-parallel decomposition with a as a terminal as desired.

Now suppose $\circ_{T_H, \boldsymbol{H}} = \circ_s$. By Lemma 42, the block-cut tree of H comprises a path of blocks from the block containing a to the block containing b, and each block is contained in $V(T_H)$. Let $B_0, \ldots, B_{k-1}$ be the blocks in order of increasing distance from a, and let $v_0, \ldots, v_k$ be defined as in Lemma 42. We have $(B_i, v_i, v_{i+1}) \in V(T_H)$ for each i, and so they have simple induced series-parallel decompositions. By Lemma 47, (B_i, v_{i+1}, v_i) also each admit simple series-parallel decompositions. Now, note that

$$(G, a, v_1) = (B_0, a, v_1) \circ_p ((K, a, b) \circ_s (B_{k-1}, v_k, v_{k-1}) \circ_s \ldots \circ_s (B_1, v_2, v_1)),$$

and each graph with terminals admits a simple series-parallel decomposition. Since B_0 has strictly fewer vertices than H, the induction hypothesis implies that G has a simple series-parallel decomposition with a as a terminal as desired. ∎

Lemma 49. *If T is a GSP decomposition of a graph G, and $(H, a, b) \in T$, then $\{a, b\}$ separates $V(H) \setminus \{a, b\}$ from the rest of G.*

Proof. We proceed by induction on $|V(G)| - |V(H)|$. If $H = G$, the claim is trivial. Otherwise, H is a proper subgraph of G and thus $\boldsymbol{H}$ is a child of some $\boldsymbol{K}$ in T. When $K = G$, the claim immediately follows from the definition of the operations $\circ_s$, $\circ_p$, $\circ_b$, $\circ_{b'}$. The remaining case is when K is also a proper subgraph of G. By the induction hypothesis, substituting K for G, the only vertices of H which could have neighbors in $V(K) \setminus V(H)$ are $\{a, b\}$. If we instead substitute K for H, we get that the only vertices of K which could have neighbors in $V(G) \setminus V(K)$ are its terminals, and whichever of these terminals are contained in H must themselves be terminals of H (one can easily verify this by considering each possible choice of $\circ_{T, \boldsymbol{K}}$). These two combined show the claim, completing the proof. ∎

We say that graphs with terminals $\boldsymbol{G} = (G, a, b)$ and $\boldsymbol{H} = (H, c, d)$ are **non-terminally disjoint** if $V(G) \cap V(H) \subseteq \{a, b, c, d\}$.

Proposition 50. *Suppose G is a 2-connected graph that does not contain any member of $\mathcal{F}_3$ as a subgraph. Let T_0, T_1 be series-parallel decompositions of G. If $\boldsymbol{H_0} \in V(T_0)$, $\boldsymbol{H_1} \in V(T_1)$ are such that both their induced decompositions are complex, then they are not non-terminally disjoint.*

Proof. Suppose otherwise, so H_0 and H_1 are non-terminally disjoint. By Corollary 46, the induced decompositions of H_0 and H_1 (in T_0 and T_1 respectively) contain (K_0, a_0, b_0) and (K_1, a_1, b_1) which themselves contain two disjoint bipaths between their respective terminals.

Since G is 2-connected, the sets $V(K_0)$ and $V(K_1)$ do not have a vertex separator of size 1, so by Menger's theorem [10, Theorem 3.3.1], there exist two disjoint paths P_0, P_1 between them. Since $\boldsymbol{K_0}$ and $\boldsymbol{K_1}$ are non-terminally disjoint, by Lemma 49, we may assume that P_0 and P_1 intersect K_0 and K_1 only at the terminals. Then the union of P_0, P_1, K_0 and K_1 contains as a subgraph an element of $\mathcal{F}_3$, a contradiction. ∎

If T is a series-parallel decomposition of a graph G, then for any two graphs $\boldsymbol{H_0}$, $\boldsymbol{H_1} \in V(T)$, either they are non-terminally disjoint or one of them contains the other as a subgraph. Therefore, Proposition 50 implies that, given any complex series-parallel decomposition T of a graph G without any member of $\mathcal{F}_3$ as a subgraph, there is a unique inclusion-minimal $\boldsymbol{H} \in V(T)$ whose induced decomposition is complex. We call this $\boldsymbol{H}$ the **minimal complex descendant** of G in T.

Lemma 51. *If $\boldsymbol{G} = (G, a, b)$ admits a series-parallel decomposition T, then $\boldsymbol{G}$ can be obtained by combining, via $\circ_p$, the graphs with terminals $(H, a, b) \in V(T)$ such that no child of (H, a, b) in T has terminals (a, b).*

Proof. Call the graphs with terminals $(H, a, b) \in V(T)$ such that no child of (H, a, b) in T has terminals (a, b) *good graphs*. We argue by induction on the number of good graphs. If there is a single good graph, then (G, a, b) itself must be good, as each graph with terminals (a, b) is either good or has two children with terminals (a, b). Otherwise, $\boldsymbol{G} = (H, a, b) \circ_p (K, a, b)$ for some subgraphs H, K. Applying the induction hypothesis on H and K and their induced decompositions shows that $\boldsymbol{H}$ and $\boldsymbol{K}$ can each be obtained by combining some subset of the good graphs via $\circ_p$, and hence so can $\boldsymbol{G}$, proving the lemma. ∎

Proposition 52. *Suppose G is a 2-connected graph without any member of $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$ as a subgraph, and let T be a series-parallel decomposition of G. Then either T is simple or G admits a simple series-parallel decomposition sharing a terminal with the minimal complex descendant of G in T.*

Proof. We may operate under the assumption that T is not simple. Let (H, a, b) be the minimal complex descendant of G in T.

Claim 52.a. *The graph with terminals (G, a, b) admits a series-parallel decomposition T'.*

▷   If $V(H) = V(G)$, we can take $T' = T$. Otherwise, by Lemma 49, $\{a, b\}$ separates $V(H) \setminus \{a, b\}$ from $V(G) \setminus V(H)$, so we are done by Proposition 41. ◀

Let T' be a series-parallel decomposition of (G, a, b).

Claim 52.b. *Suppose $(K, a, b) \in V(T')$ has no child in T' with terminals (a, b). Then K is a subgraph of either H or $G - (H - \{a, b\})$.*

$\triangleright$ If K is a single edge, then it is a subgraph of H or of $G - (H - \{a, b\})$ depending on whether a and b are neighbors in H. Otherwise, the induced decomposition on $\boldsymbol{K}$ in T' is nontrivial, and we must have $\circ_{T', \boldsymbol{K}} = \circ_s$ as otherwise $\boldsymbol{K}$ would have a child with terminals (a, b). By Lemma 42, the block-cut graph of K comprises a path, with the blocks containing a and b as distinct endpoints. If K contains vertices of both $H - \{a, b\}$ and $G - H$, one of these blocks, say B, must contain vertices from both. By Lemma 49 applied to $\boldsymbol{H}$ with respect to T, $\{a, b\}$ must separate $H - \{a, b\}$ from $G - H$, but B cannot contain both a and b, so it remains connected after the removal of $\{a, b\}$, a contradiction. $\blacktriangleleft$

By Lemma 51, we can write

$$(G, a, b) = (K_1, a, b) \circ_p \cdots \circ_p (K_\ell, a, b),$$

where $\boldsymbol{K_1}, \ldots, \boldsymbol{K_\ell} \in V(T')$ are such that their children do not have terminals (a, b). We may order these graphs so that for some $1 \leqslant m \leqslant \ell$, the K_i with $i < m$ are subgraphs of $G - (H - \{a, b\})$ and the K_i with $i \geqslant m$ are subgraphs of H. Then $\bigcup_{i=m}^{\ell} K_i = H$, so

$$(G, a, b) = (K_1, a, b) \circ_p \ldots \circ_p (K_{m-1}, a, b) \circ_p (H, a, b).$$

Since (H, a, b) is the minimal complex descendant of G in T, the induced decompositions of its children are simple and $\circ_{T, \boldsymbol{H}} = \circ_p$. Let (H_0, a, b) and (H_1, a, b) be its children in T. Then

$$(G, a, b) = ((K_1, a, b) \circ_p \cdots \circ_p (K_{m-1}, a, b) \circ_p (H_0, a, b)) \circ_p (H_1, a, b).$$

Consider any K_i with $i < m$. Since it is a subgraph of $G - (H - \{a, b\})$, its induced decomposition in T' cannot be complex, as otherwise H and K_i contradict Proposition 50. Furthermore, (K_i, a, b) must be bridged, as otherwise, by Lemma 45, it would contain a bipath between a and b which is internally disjoint from the two bipaths between a and b contained in H (also given by Lemma 45), and so G would contain as a subgraph a member of $\mathcal{F}_2$.

Hence each (K_i, a, b) is bridged and has a simple induced decomposition in T' of complexity 0. By connecting each of these decompositions with the induced decomposition of (H_0, a, b) from T, we obtain a simple series-parallel decomposition of $(K_1, a, b) \circ_p \ldots \circ_p (K_{m-1}, a, b) \circ_p (H_0, a, b)$. Finally, we may apply Proposition 48 to show that G admits a simple series-parallel decomposition with a as a terminal. $\blacksquare$

8.3 The general case

In this subsection we complete the proof of implication (2) $\implies$ (3) in Theorem 7 by handling non-2-connected graphs.

Lemma 53. *Let G be a connected graph which does not contain any member of $\mathcal{F}_3$ as a subgraph. Suppose H, K are distinct leaf blocks of G with respective cut-vertices v_H, v_K. If T_H, T_K are series-parallel decompositions of H and K respectively, and $(H', a, b) \in T_H$ has a complex induced decomposition with $v_H \notin V(H')$, then any $\boldsymbol{K'} \in T_K$ with a complex induced decomposition must have $v_K \in V(K')$.*

Proof. By 2-connectedness of H, there exists a path from a to v_H in $G-\{b\}$. By Lemma 49, any such path of minimum length cannot contain any non-terminal vertex of H'. Likewise, we can find a path from b to v_H which does not contain a non-terminal vertex of H'.

If K' does not contain v_K, we can similarly find paths from v_K to each of its terminals which does not contain the non-terminal vertices of K'. Finally, we can take any path from v_H to v_K to construct paths between each pair of terminals of H' and K' each avoiding the non-terminal vertices of H' and K'. Since each of H', K' contain a pair of bipaths between their terminals (Lemma 45), G will contain as a subgraph a member of $\mathcal{F}_3$, contradicting the assumption. ∎

Theorem 54. *Suppose G is a connected graph which does not contain any member of $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$ as a subgraph. Then G has a simple GSP decomposition.*

Proof. We proceed by induction on the number of blocks in G. The base case when G has only one block is handled by Proposition 52. Now suppose G has at least 2 blocks and let H, K be distinct leaf blocks with respective cut-vertices v_H, v_K. By Proposition 41, we can find series-parallel decompositions T_H, T_H of H and K respectively with each having v_H, v_K respectively as a terminal (we can make the terminals be v_H/v_K and a neighbor, for instance).

If at least one of T_H, T_K is simple, then by the inductive hypothesis we can find a simple decomposition of the rest of the graph, and then, by Lemma 43, we get a simple decomposition of G. Otherwise, both T_H and T_K are complex, so Lemma 53 implies that at least one (without loss of generality, let it be H) has its minimal complex descendant contain v_H. Since $\boldsymbol{H}$ has v_H as a terminal, v_H must be a terminal of this minimal complex descendant. Hence by Proposition 52, we can find a simple decomposition of H with v_H as a terminal, then again we can find a simple decomposition of the rest of the graph, and by Lemma 43, we have a simple decomposition of G. ∎

9 Searching subdivisions

9.1 Introduction

In this section we prove implication (3) $\implies$ (1) of Theorem 7; i.e., we show that every graph G with a simple GSP decomposition has topological inspection number at most 3. We begin by noting that if a graph G has a simple GSP decomposition, then so do all its subdivisions:

Lemma 55. *Let $\boldsymbol{G} = (G, a, b)$ be a graph with terminals and let T be a GSP decomposition of $\boldsymbol{G}$. Suppose that G' is a subdivision of G. For each subgraph D of G, let D' denote the corresponding subgraph of G' obtained by subdividing D.*

Then $\boldsymbol{G'} = (G, a, b)$ has a GSP decomposition T' such that for all $\boldsymbol{D} = (D, s, t) \in V(T)$, we have $\boldsymbol{D'} = (D', s, t) \in V(T')$ and $\circ_{T, \boldsymbol{D}} = \circ_{T', \boldsymbol{D'}}$. Additionally, if T is simple, then T' is simple too.

Proof. Replace each node (D, s, t) in T with (D', s, t). The leaf nodes now correspond to subdivided edges (i.e., paths), which have series-parallel decompositions using only the series operation. Replacing each leaf node with its series-parallel decomposition gives the desired GSP decomposition T'. Note that the series-parallel decompositions replacing leaf nodes have complexity 0. Since whether a graph with terminals is bridged is not changed by subdivision, the complexity of each subtree in the decomposition remains the same as before, and hence if T is simple, then T' is simple as well. ∎

In view of Lemma 55, to prove implication (3) $\implies$ (1) of Theorem 7, we just need to argue that every graph G with a simple GSP decomposition T has a subdivision G' such that $\mathsf{in}(G') \leqslant 3$. Indeed, by Lemma 55, this would imply that every subdivision of G has a further subdivision with inspection number at most 3, as desired.

9.2 Aligned searches and quotient graphs

Given a graph G with a simple GSP decomposition T, we shall construct a subdivision G' of G with $\mathsf{in}(G') \leqslant 3$ inductively by "amalgamating" successful 3-searches on subdivisions of the descendants of G in T. In order for this amalgamation process to work, we need these 3-searches to satisfy a certain property defined as follows:

Definition 56 (Aligned Searches). A k-search $\mathcal{S}$ of length ℓ on a graph G is **aligned** to an ordered pair of vertices $(a, b) \in V(G)^2$ if $a \in \mathsf{PC}_t(\mathcal{S})$ for each $1 \leqslant t \leqslant \ell$, and $b \notin \mathsf{FC}_t(\mathcal{S})$ for each $0 \leqslant t < \ell$.

If $A \subseteq V(G)$ and $\mathcal{S}$ is a k-search of length ℓ, we say that $\mathcal{S}$ is **aligned to** $(a, b) \in V(G)^2$ **over** A if $a \in \mathsf{PC}_t(\mathcal{S}, A)$ for each $1 \leqslant t \leqslant \ell$, and $b \notin \mathsf{FC}_i(\mathcal{S}, A)$ for each $0 \leqslant t < \ell$.

Here $\mathsf{PC}_t(\mathcal{S}, A)$ and $\mathsf{FC}_t(\mathcal{S}, A)$ refer to the sets of pre-cleared and fully cleared vertices if the set A is initially cleared (see Definition 9).

The next definition provides the terminology and notation that will be used in describing the amalgamation process:

Definition 57 (Quotient Graphs). Let G be a graph and let $\mathcal{E}$ be an equivalence relation on $V(G)$. The **quotient graph** $G/\mathcal{E}$ is the graph with vertex set $V(G)/\mathcal{E}$ that includes an edge between distinct $\mathcal{E}$-equivalence classes $c_1, c_2 \in V(G)/\mathcal{E}$ if and only if there are adjacent vertices $v_1 \in c_1$ and $v_2 \in c_2$.

A subset $X \subseteq V(G)$ is $\mathcal{E}$-**invariant** if it is a union of $\mathcal{E}$-equivalence classes. For an $\mathcal{E}$-invariant set $X \subseteq V(G)$, we write $X/\mathcal{E}$ to denote the set of all $\mathcal{E}$-equivalence classes contained in X.

Define two maps $\wedge_{\mathcal{E}}, \vee_{\mathcal{E}} \colon \mathcal{P}(V(G)) \to \mathcal{P}(V(G)/\mathcal{E})$ via

$$\wedge_{\mathcal{E}}(X) = \{c \in V(G)/\mathcal{E} \, : \, c \subseteq X\} \quad \text{and} \quad \vee_{\mathcal{E}}(X) = \{c \in V(G)/\mathcal{E} \, : \, c \cap X \neq \varnothing\}.$$

Note that if X is $\mathcal{E}$-invariant, then $\wedge_{\mathcal{E}}(X) = \vee_{\mathcal{E}}(X) = X/\mathcal{E}$.

Let $\mathcal{S}$ be a k-search of length ℓ on G and let $A \subseteq V(G)$. We say that $\mathcal{S}$ is $\mathcal{E}$-**invariant over** A if for all $1 \leqslant t \leqslant \ell$, the set $\mathsf{PC}_t(\mathcal{S}, A)$ is $\mathcal{E}$-invariant. Given a k-search $\mathcal{S} = (S_t)_{t=1}^{\ell}$ on G, we let $\vee_{\mathcal{E}}(\mathcal{S}) = (\vee_{\mathcal{E}}(S_t))_{t=1}^{\ell}$.

If $\mathcal{E}$ is an equivalence relation on $V(G)$ and $\mathcal{S}$ is a k-search on G, then $\vee_{\mathcal{E}}(\mathcal{S})$ is a k-search on $G/\mathcal{E}$. The next proposition allows us to compute the sets of pre-cleared and fully cleared vertices for it.

Proposition 58. *Let G be a graph and let $\mathcal{E}$ be an equivalence relation on $V(G)$. Suppose that $A \subseteq V(G)$ and $\mathcal{S}$ is a k-search of length ℓ on G that is $\mathcal{E}$-invariant over A. Then:*

- $\mathsf{FC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(\mathsf{FC}_t(\mathcal{S}, A))$ *for all $0 \leqslant t \leqslant \ell$;*

- $\mathsf{PC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \mathsf{PC}_t(\mathcal{S}, A)/\mathcal{E}$ *for all $1 \leqslant t \leqslant \ell$.*

Proof. We shall prove the following two implications for all $1 \leqslant t \leqslant \ell$:

(a) If $\mathsf{FC}_{t-1}(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(\mathsf{FC}_{t-1}(\mathcal{S}, A))$, then $\mathsf{PC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \mathsf{PC}_t(\mathcal{S}, A)/\mathcal{E}$.

(b) If $\mathsf{PC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \mathsf{PC}_t(\mathcal{S}, A)/\mathcal{E}$, then $\mathsf{FC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(\mathsf{FC}_t(\mathcal{S}, A))$.

Together with the observation that $\mathsf{FC}_0(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(A) = \wedge_{\mathcal{E}}(\mathsf{FC}_0(\mathcal{S}, A))$, these immediately yield the desired result.

First suppose that $\mathsf{FC}_{t-1}(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(\mathsf{FC}_{t-1}(\mathcal{S}, A))$. Then, by definition,

$$\mathsf{PC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(\mathsf{FC}_{t-1}(\mathcal{S}, A)) \cup \vee_{\mathcal{E}}(S_t) \quad \text{and} \quad \mathsf{PC}_t(\mathcal{S}, A) = \mathsf{FC}_{t-1}(\mathcal{S}, A) \cup S_t.$$

Since the set $\mathsf{PC}_t(\mathcal{S}, A)$ is $\mathcal{E}$-invariant, every $\mathcal{E}$-equivalence class c is either contained in it or disjoint from it. Therefore,

$$\begin{aligned} c \in \mathsf{PC}_t(\mathcal{S}, A)/\mathcal{E} &\iff c \subseteq \mathsf{FC}_{t-1}(\mathcal{S}, A) \text{ or } c \cap S_t \neq \varnothing \\ &\iff c \in \wedge_{\mathcal{E}}(\mathsf{FC}_{t-1}(\mathcal{S}, A)) \text{ or } c \in \vee_{\mathcal{E}}(S_t) \\ &\iff c \in \mathsf{PC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)). \end{aligned}$$

Now suppose that $\mathsf{PC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \mathsf{PC}_t(\mathcal{S}, A)/\mathcal{E}$. Then $\mathsf{FC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A))$ comprises exactly those equivalence classes $c \in \mathsf{PC}_t(\mathcal{S}, A)/\mathcal{E}$ that satisfy $c \cap \partial_G(\mathsf{PC}_t(\mathcal{S}, A)) = \varnothing$, i.e., that are fully contained in $\mathsf{FC}_t(\mathcal{S}, A)$. Therefore, $\mathsf{FC}_t(\vee_{\mathcal{E}}(\mathcal{S}), \wedge_{\mathcal{E}}(A)) = \wedge_{\mathcal{E}}(\mathsf{FC}_t(\mathcal{S}, A))$, as desired. ∎

9.3 Clearing balls

Recall that for $v \in V(G)$ and $r \in \mathbb{N}$, $B_G(v, r)$ denotes the ball of radius r around v in G. In this subsection we establish two useful auxiliary constructions that allow us to clear a ball of large radius around a vertex in a subdivision of a given graph.

Lemma 59. *Let G be a graph and let $v, w \in V(G)$ be distinct vertices. Let d be the degree of w in G. Fix some $r \in \mathbb{N}$ and let H be a subdivision of G such that each edge incident to w is subdivided at least $2^{d-1}r + 1$ times. Then there exists a 3-search $\mathcal{S}$ on H of length $\ell = (2^d - 1)r$ such that $w \in S_t$ for all $1 \leqslant t \leqslant \ell$, $\mathcal{S}$ is aligned to (w, v), and $\mathsf{FC}_*(\mathcal{S}) \supseteq B_H(w, r)$.*

Proof. Let $P_1, \ldots, P_d$ be the paths in H replacing the edges of G incident to w. Denote the endpoint of P_i distinct from w by u_i. Let H' be the graph obtained from H by removing all the vertices in $V(P_1) \cup \ldots \cup V(P_d)$ except for $u_1, \ldots, u_d$, and let H^* be the graph obtained as a disjoint union of H' and d paths $P_1^*, \ldots, P_d^*$, where each P_i^* is of the same length as P_i. Denote the two endpoints of P_i^* by w_i and u_i^* and define an equivalence relation $\mathcal{E}$ on $V(H^*)$ by making the following sets the only nontrivial equivalence classes: $\{w_1, \ldots, w_d\}, \{u_1, u_1^*\}, \ldots, \{u_d, u_d^*\}$. Then there is an obvious isomorphism $H \cong H^*/\mathcal{E}$, which allows us to identify H with $H^*/\mathcal{E}$.

For $1 \leqslant i \leqslant d$, denote by $V_i[a, b]$ the set of all vertices in P_i^* at distance between a and b (inclusive) from w_i. Let $\mathcal{S}_i$ be the $(d+2)$-search on H^* of length $2^{d-i}r$, where for each $1 \leqslant t \leqslant 2^{d-i}r$,

$$(S_i)_t = \{w_1, \ldots, w_d\} \cup V_i[t, t+1].$$

Note that since the length of P_i^* is at least $2^{d-1}r + 2$, $u_i^* \notin (S_i)_t$. Let $\mathcal{S}^* = \mathcal{S}_1 ^\frown \cdots ^\frown \mathcal{S}_d$, where $^\frown$ denotes concatenation of sequences, and define $\mathcal{S} = \vee_\mathcal{E}(\mathcal{S}^*)$. We claim that $\mathcal{S}$ is as desired.

It is clear that $\mathcal{S}$ is a 3-search (since the vertices $w_1, \ldots, w_d$ represent the same equivalence class). The length of $\mathcal{S}$ is equal to the length of $\mathcal{S}^*$, which is

$$2^{d-1}r + 2^{d-2}r + \cdots + r = (2^d - 1)r.$$

Also, for each t, w is the unique vertex of G in S_t. This implies that $\mathcal{S}$ is aligned to (w, v).

Next we observe that the search $\mathcal{S}^*$ is $\mathcal{E}$-invariant. Indeed, the equivalence class $\{w_1, \ldots, w_d\}$ is contained in every S_t^*, so all its elements are always pre-cleared. On the other hand, the vertices in an equivalence class of the form $\{u_i, u_i^*\}$ are never pre-cleared.

It is straightforward to verify that $V_i[0, 2^{d-i}r] \subseteq \mathsf{FC}_*(\mathcal{S}_i)$ and hence $V_i[0, 2^{d-i}r] \subseteq \mathsf{FC}_{(2^d - 2^{d-i})r}(\mathcal{S}^*)$. It follows that for each $t \geqslant 1$,

$$V_i[0, 2^{d-i}r - t + 1] \subseteq \mathsf{PC}_{(2^d - 2^{d-i})r + t}(\mathcal{S}^*) \quad \text{and} \quad V_i[0, 2^{d-i}r - t] \subseteq \mathsf{FC}_{(2^d - 2^{d-i})r + t}(\mathcal{S}^*).$$

By applying this with $t = (2^{d-i} - 1)r$, we get that $V_i[0, r] \subseteq \mathsf{FC}_*(\mathcal{S}^*)$ for each i. Hence, by Proposition 58, $\mathsf{FC}_*(\mathcal{S})$ contains all the vertices of H at distance at most r from w, as desired. ∎

Lemma 60. *Let G be a graph and $v, w \in V(G)$ be distinct vertices. Let d be the degree of v in G. Fix some $r \in \mathbb{N}$ and let H be a subdivision of G such that each edge incident to v is subdivided at least $2^{d-1}r$ times. Let $A = V(H) \setminus B_H(v, r)$. Then there exists a 3-search $\mathcal{S}$ on H of length $\ell = (2^d - 1)r$ such that $v \in S_t$ for all $1 \leqslant t \leqslant \ell$, $\mathcal{S}$ is aligned to (w, v) over A, and $\mathsf{FC}_*(\mathcal{S}, A) = V(H)$.*

Proof. Let $P_1, \ldots, P_d$ be the paths in H replacing the edges of G incident to v. Denote the endpoint of P_i distinct from v by u_i. Let H' be the graph obtained from H by removing all the vertices in $V(P_1) \cup \ldots \cup V(P_d)$ except for $u_1, \ldots, u_d$, and let H^* be the graph obtained as a disjoint union of H' and d paths $P_1^*, \ldots, P_d^*$, where each P_i^* is of the same length as P_i. Denote the two endpoints of P_i^* by v_i and u_i^* and define an equivalence

relation $\mathcal{E}$ on $V(H^*)$ by making the following sets the only nontrivial equivalence classes: $\{v_1, \ldots, v_d\}$, $\{u_1, u_1^*\}$, $\ldots$, $\{u_d, u_d^*\}$. Then there is an obvious isomorphism $H \cong H^*/\mathcal{E}$, which allows us to identify H with $H^*/\mathcal{E}$. Let $A^* \subseteq V(H^*)$ be the $\mathcal{E}$-invariant set such that $A^*/\mathcal{E} = A$. Explicitly, A^* comprises $V(H')$ together with the vertices on each of the paths P_i at distance more than r from v_i.

For $1 \leqslant i \leqslant d$, denote by $V_i[a, b]$ the set of all vertices in P_i^* at distance between a and b (inclusive) from v_i. Let $\mathcal{S}_i$ be the $(d+2)$-search on H^* of length $2^{i-1}r$ where for each $1 \leqslant t \leqslant 2^{i-1}r$,

$$(S_i)_t = \{v_1, \ldots, v_d\} \cup V_i[2^{i-1}r - t, 2^{i-1}r - t + 1].$$

Let $\mathcal{S}^* = \mathcal{S}_1 \frown \cdots \frown \mathcal{S}_d$ and $\mathcal{S} = \vee_{\mathcal{E}}(\mathcal{S}^*)$. We claim that $\mathcal{S}$ is as desired.

It is clear that $\mathcal{S}$ is a 3-search (since the vertices v_1, $\ldots$, v_d represent the same equivalence class). The length of $\mathcal{S}$ is equal to the length of $\mathcal{S}^*$, which is

$$r + 2r + 2^2 r + \cdots + 2^{d-1} r = (2^d - 1)r.$$

By construction, $v \in S_t$ for all t.

For $t < (2^{i-1} - 1)r$ (the first turn of $\mathcal{S}_i$ in $\mathcal{S}'$), we have $V(P_i) \setminus \mathsf{FC}_t(\mathcal{S}^*, A^*) = V_i[0, r + t]$. Thus,

$$V(P_i) \setminus \mathsf{FC}_{(2^{i-1}-1)r}(\mathcal{S}^*, A^*) = V_i[0, 2^{i-1}r].$$

By turn $(2^i - 1)r$, all the vertices in P_i will become fully cleared and will remain so for the remainder of $\mathcal{S}^*$. Since $V(H') \subseteq A^*$, we conclude that $\mathsf{FC}_*(\mathcal{S}^*, A^*) = V(H^*)$.

Note that each v_i is in S_t^* for all t, while u_i and u_i^* are in $\mathsf{PC}_t(\mathcal{S}^*, A^*)$ for all t. This shows that $\mathcal{S}^*$ is $\mathcal{E}$-invariant over A^*. Hence, by Proposition 58, $\mathsf{FC}_*(\mathcal{S}, A) = V(H)$, as desired. Finally, $\mathcal{S}$ is aligned to (w, v) over A, since every vertex of G except v is fully cleared throughout, while v does not become fully cleared until the last move. ∎

9.4 Amalgamation

In this subsection we combine aligned searches by the GSP operations.

Proposition 61. *Let (G_0, a, c) and (G_1, c, b) be connected graphs with terminals such that $V(G_0) \cap V(G_1) = \{c\}$ and let $(G, a, b) = (G_0, a, c) \circ_s (G_1, c, b)$. Suppose that for each $i \in \{0, 1\}$, there is a subdivision H_i of G_i admitting a successful 3-search $\mathcal{S}_i$ aligned to the corresponding terminals. Then there is a subdivision of G admitting a successful 3-search aligned to (a, b).*

Proof. The proof is illustrated in Fig. 11. The desired subdivision of G is simply the union of H_0 and H_1, and the desired 3-search is the concatenation $\mathcal{S}_0 \frown \mathcal{S}_1$. To aid in our analysis, we let H be the *disjoint* union of the graphs H_0 and H_1 and denote the copy of c in each $H_i \subset H$ by c_i. Let $\mathcal{E}$ be the equivalence relation on $V(H)$ with $\{c_0, c_1\}$ as the only nontrivial equivalence class. Then $H/\mathcal{E}$ can be naturally identified with the (non-disjoint) union of H_0 and H_1.
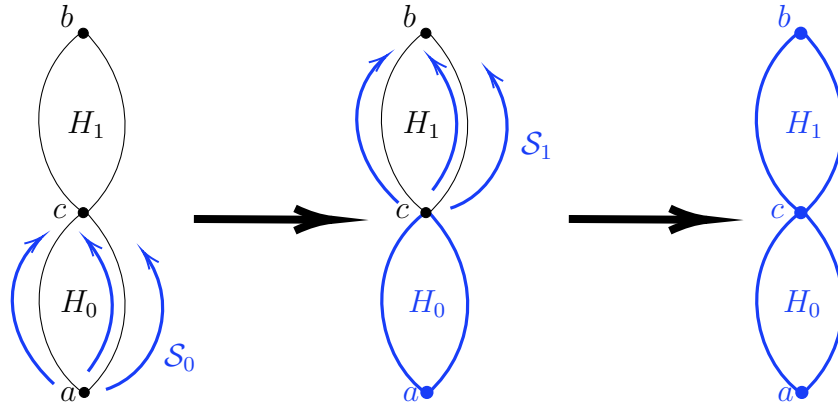
Figure 11: A pictorial representation of the construction in Proposition 61.

Construct searches $\mathcal{S}_0'$ and $\mathcal{S}_1'$ on H from $\mathcal{S}_0$ and $\mathcal{S}_1$ respectively by replacing every appearance of c by both c_0 and c_1, and let $\mathcal{S} = \mathcal{S}_0' {}^\frown \mathcal{S}_1'$. Then $\vee_\mathcal{E}(\mathcal{S})$ is a 3-search on $H/\mathcal{E}$. (If we identify $H/\mathcal{E}$ with the union of H_0 and H_1, then $\vee_\mathcal{E}(\mathcal{S}) = \mathcal{S}_0 {}^\frown \mathcal{S}_1$.)

We now check that $\mathcal{S}$ is $\mathcal{E}$-invariant. To this end, we need to show that for all t, $c_0 \in \mathsf{PC}_t(\mathcal{S}) \iff c_1 \in \mathsf{PC}_t(\mathcal{S})$. We claim that for $t \leqslant |\mathcal{S}_0|$,

$$c_0 \in \mathsf{PC}_t(\mathcal{S}) \iff c_0 \in S_t \iff c_1 \in S_t \iff c_1 \in \mathsf{PC}_t(\mathcal{S}).$$

The first equivalence follows since $\mathcal{S}_0$ is aligned to (a, c) and hence $c_0 \notin \mathsf{FC}_{t-1}(\mathcal{S})$ for any $t \leqslant |\mathcal{S}_0|$. The second equivalence holds by the construction of $\mathcal{S}$. The final equivalence holds since no neighbor of c_1 is pre-cleared during the first $|\mathcal{S}_0|$ moves.

Now suppose $t > |\mathcal{S}_0|$. Then both c_0 and c_1 are pre-cleared. Indeed, c_0 is fully cleared since $\mathcal{S}_0$ is a successful search on H_0, while c_1 is pre-cleared since $\mathcal{S}_1$ is aligned to (c, b).

Since $\mathcal{S}$ is $\mathcal{E}$-invariant, we may apply Proposition 58 to conclude that $\vee_\mathcal{E}(\mathcal{S})$ is successful and aligned to (a, b), as desired. ∎

Proposition 62. *Let (G_0, a, c) and (G_1, a, c) be connected graphs with terminals such that $V(G_0) \cap V(G_1) = \{a\}$ and let $(G, a, b) = (G_0, a, b) \circ_b (G_1, a, c)$. Suppose that there is a subdivision H_1 of G_1 with a successful 3-search $\mathcal{S}_1$ aligned to (a, c), and every subdivision of G_0 has a further subdivision with a successful 3-search aligned to (a, b). Then there is a subdivision of G admitting a successful 3-search aligned to (a, b).*

Proof. The proof is illustrated in Fig. 12. Let d be the degree of a in G_0 and let H_0 be a subdivision of G_0 such that every edge incident to a is subdivided at least $2^{d-1} |\mathcal{S}_1| + 1$ times and there is a successful 3-search $\mathcal{S}_0$ on H_0 aligned to (a, b). By Lemma 59, there exists a 3-search $\mathcal{S}_a$ on H_0 aligned to (a, b) with $a \in (S_a)_t$ for all t such that $B_{H_0}(a, |\mathcal{S}_1|) \subseteq \mathsf{FC}_*(\mathcal{S}_a)$. The desired subdivision of G is the union of H_0 and H_1, and the desired 3-search is the concatenation $\mathcal{S}_a {}^\frown \mathcal{S}_1 {}^\frown \mathcal{S}_0$.

Let H be the disjoint union of H_0 and H_1 and denote the copy of a in each $H_i \subset H$ by a_i. Let $\mathcal{E}$ be the equivalence relation on $V(H)$ with $\{a_0, a_1\}$ as the only nontrivial

equivalence class. Then $H/\mathcal{E}$ can be naturally identified with the (non-disjoint) union of H_0 and H_1.

Construct searches $\mathcal{S}_a'$, $\mathcal{S}_0'$, and $\mathcal{S}_1'$ on H from $\mathcal{S}_a$, $\mathcal{S}_0$, and $\mathcal{S}_1$ respectively by replacing every appearance of a with both a_0 and a_1. Let $\mathcal{S} = \mathcal{S}_a' {}^\frown \mathcal{S}_1' {}^\frown \mathcal{S}_0'$. For brevity, let $t_0 = |\mathcal{S}_a|$ and $t_1 = t_0 + |\mathcal{S}_1|$.

We now check that $\mathcal{S}$ is $\mathcal{E}$-invariant. To this end, we need to show that for all t, $a_0 \in \mathsf{PC}_t(\mathcal{S}) \iff a_1 \in \mathsf{PC}_t(\mathcal{S})$. By construction, $a_0 \in S_t$ for $t \leqslant t_0$. Since $B_H(a_0, |\mathcal{S}_1|) \subseteq \mathsf{FC}_{t_0}(\mathcal{S})$, we have $a_0 \in \mathsf{PC}_t(\mathcal{S})$ for $t_0 < t \leqslant t_1$. Since $\mathcal{S}_0$ is aligned to (a, b), $a_0 \in \mathsf{PC}_t(\mathcal{S})$ for $t \geqslant t_1$. Thus, $a_0 \in \mathsf{PC}_t(\mathcal{S})$ for every t.

For $t \leqslant t_0$, we have $a_1 \in S_t$ by construction. Since $\mathcal{S}_1$ is aligned to (a, c), $a_1 \in \mathsf{PC}_t(\mathcal{S})$ for $t_0 < t \leqslant t_1$. Since $\mathcal{S}_1$ is successful, $a_1 \in \mathsf{FC}_t(\mathcal{S})$ for $t > t_1$. Thus $a_1 \in \mathsf{PC}_t(\mathcal{S})$ for every t too.

We have thus shown that $\mathcal{S}$ is $\mathcal{E}$-invariant. Note that $\vee_\mathcal{E}(\mathcal{S})$ is a 3-search on $H/\mathcal{E}$ (if we identify $H/\mathcal{E}$ with the union of H_0 and H_1, then $\vee_\mathcal{E}(\mathcal{S}) = \mathcal{S}_a {}^\frown \mathcal{S}_1 {}^\frown \mathcal{S}_0$), and furthermore, by Proposition 58, $\vee_\mathcal{E}(\mathcal{S})$ is successful and aligned to (a, b), as desired. ∎

Proposition 63. *Let (G_0, a, b) and (G_1, b, c) be connected graphs with terminals such that $V(G_0) \cap V(G_1) = \{b\}$ and let $(G, a, b) = (G_0, a, b) \circ_{b'} (G_1, b, c)$. Suppose that there is a subdivision H_1 of G_1 with a successful 3-search $\mathcal{S}_1$ aligned to (c, b), and every subdivision of G_0 has a further subdivision with a successful 3-search aligned to (a, b). Then there is a subdivision of G admitting a successful 3-search aligned to (a, b).*

Proof. The proof is illustrated in Fig. 13. Let d be the degree of b in G_0 and let H_0 be a subdivision of G_0 such that every edge incident to b is subdivided at least $2^{d-1}(|\mathcal{S}_1| + 1)$ times and there is a successful 3-search $\mathcal{S}_0$ on H_0 aligned to (a, b). For our analysis, it will be convenient to modify $\mathcal{S}_0$ so that the vertex b remains uncleared even after the last step. To this end, we let $\mathcal{S}_0^*$ be the search on H_0 obtained from $\mathcal{S}_0$ by removing b from the set $(S_0)_{|\mathcal{S}_0|}$ (i.e., from the last searched set in $\mathcal{S}_0$).

Let $A = V(H_0) \setminus B_{H_0}(b, |\mathcal{S}_1| + 1)$. By Lemma 60, there exists a 3-search $\mathcal{S}_b$ on H_0 aligned to (a, b) over A with $b \in (S_b)_t$ for all t such that $\mathsf{FC}_*(\mathcal{S}_b, A) = V(H_0)$. The desired
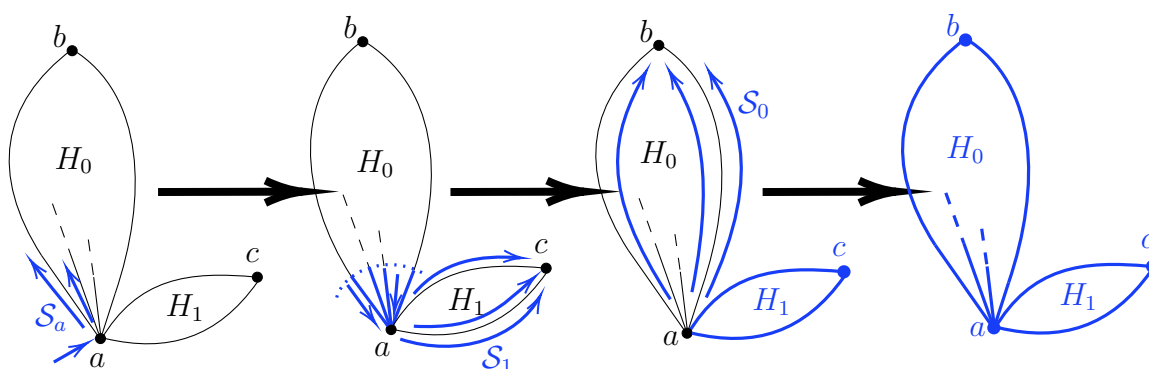

Figure 12: A pictorial representation of the construction in Proposition 62.

subdivision of G is the union of H_0 and H_1, and the desired 3-search is the concatenation $\mathcal{S}_0^{*\frown}\mathcal{S}_1^{\frown}\mathcal{S}_b$.

Let H be the disjoint union of H_0 and H_1 and denote the copy of b in each $H_i \subset H$ by b_i. Let $\mathcal{E}$ be the equivalence relation on $V(H)$ with $\{b_0, b_1\}$ as the only nontrivial equivalence class. Then $H/\mathcal{E}$ can be naturally identified with the (non-disjoint) union of H_0 and H_1.

We construct searches $\mathcal{S}_0'$, $\mathcal{S}_1'$, and $\mathcal{S}_b'$ on H from $\mathcal{S}_0^*$, $\mathcal{S}_1$, and $\mathcal{S}_b$ respectively by replacing every appearance of b with both b_0 and b_1. Let $\mathcal{S} = \mathcal{S}_0'^{\frown}\mathcal{S}_1'^{\frown}\mathcal{S}_b'$. For brevity, let $t_0 = |\mathcal{S}_0|$, and $t_1 = t_0 + |\mathcal{S}_1|$.

We now check that $\mathcal{S}$ is $\mathcal{E}$-invariant. To this end, we claim that for all t,

$$b_0 \in \mathsf{PC}_t(\mathcal{S}) \iff b_0 \in S_t \iff b_1 \in S_t \iff b_1 \in \mathsf{PC}_t(\mathcal{S}).$$

The second equivalence holds by the construction of $\mathcal{S}$, so we only need to establish the first and third equivalences. We start by considering the case $t \leqslant t_0$. The first equivalence then follows since $\mathcal{S}_0$ is aligned to (a, b) and hence $b_0 \notin \mathsf{FC}_{t-1}(\mathcal{S})$, while the third equivalence holds since no neighbor of b_1 is pre-cleared during the first t_0 moves.

Next we consider the case $t_0 < t \leqslant t_1$. Since $\mathcal{S}_0$ is aligned to (a, b) and, by construction, $b \notin (S_0^*)_{t_0}$, we have $b \notin \mathsf{PC}_{t_0}(\mathcal{S})$. Since $\mathsf{PC}_*(\mathcal{S}_0) = V(H_0)$, we have $\mathsf{PC}_{t_0}(\mathcal{S}) = V(H_0) \setminus \{b_0\}$ and hence $\mathsf{FC}_{t_0}(\mathcal{S}) = V(H_0) \setminus B_{H_0}(b_0, 1)$. Therefore, if $t_0 < t \leqslant t_1$, then no neighbor of b_0 is in $\mathsf{PC}_t(\mathcal{S})$. Thus $b_0 \notin \mathsf{FC}_t(\mathcal{S})$ for such t, which yields $b_0 \in \mathsf{PC}_t(\mathcal{S}) \iff b_0 \in S_t$. Since $\mathcal{S}_1$ is aligned to (c, b), $b_1 \notin \mathsf{FC}_t(\mathcal{S})$ for $t < t_1$, so $b_1 \in \mathsf{PC}_t(\mathcal{S}) \iff b_1 \in S_t$ as well.

Finally, if $t > t_1$, then $b_0, b_1 \in (S_b)_t \subseteq \mathsf{PC}_t(\mathcal{S})$ always. We have thus shown that the search $\mathcal{S}$ is $\mathcal{E}$-invariant. Note that $\vee_{\mathcal{E}}(\mathcal{S})$ is a 3-search on $H/\mathcal{E}$, and furthermore, by Proposition 58, it is successful and aligned to (a, b), proving the claim. $\blacksquare$

Proposition 64. *Let (G_0, a, b), (G_1, a, c), and (G_2, d, b) be connected graphs with terminals such that $V(G_0) \cap V(G_1) = \{a\}$, $V(G_0) \cap V(G_2) = \{b\}$, and $V(G_1) \cap V(G_2) = \varnothing$. Let*

$$(G, a, b) = (G_0, a, b) \circ_p ((G_1, a, c) \circ_s (e_{cd}, c, d) \circ_s (G_2, d, b)),$$

where e_{cd} is the graph comprising a single edge between c and d. Suppose that for each $i \in \{1, 2\}$, there exists a subdivision H_i of G_i with a successful 3-search $\mathcal{S}_i$ aligned to (a, c)
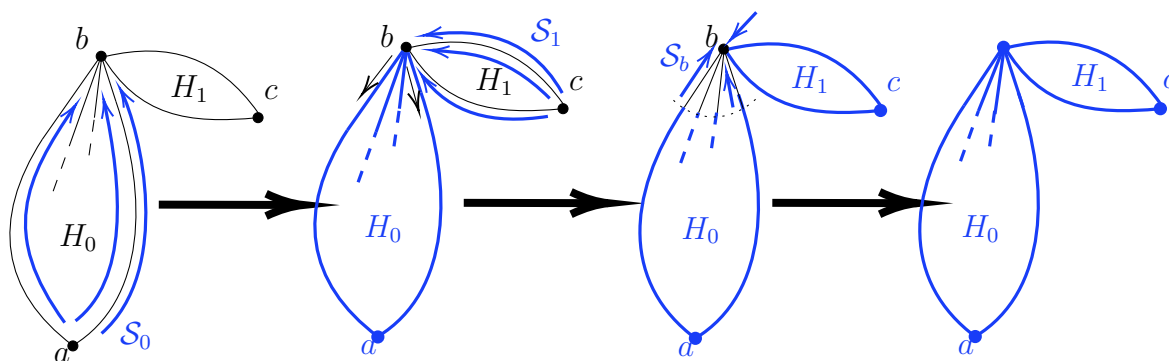

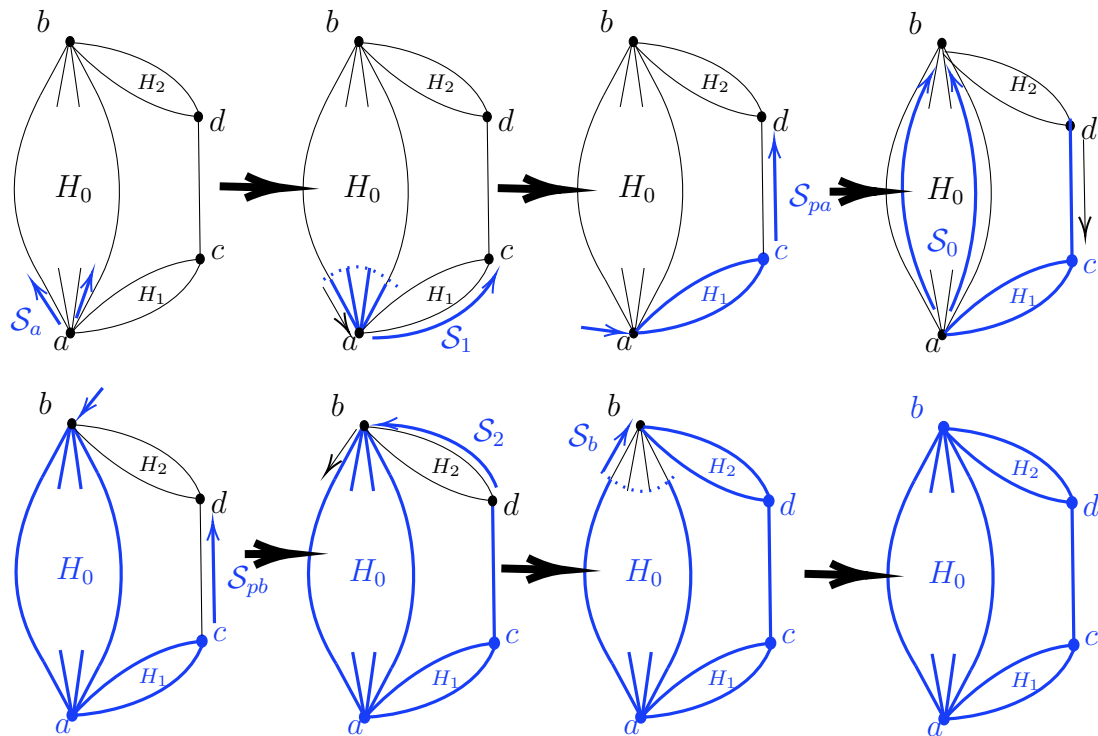**Figure 13:** A pictorial representation of the construction in Proposition 63.

Figure 14: A pictorial representation of the construction in Proposition 64.

or (d, b) respectively, and every subdivision of G_0 has a further subdivision with a successful 3-search aligned to (a, b). Then there is a subdivision of G admitting a successful 3-search aligned to (a, b).

Proof. The proof is illustrated in Fig. 14. Let Δ be the maximum of the degrees of a and b in G_0 and let H_0 be a subdivision of G_0 such that every edge incident to a is subdivided at least $2^{\Delta-1}|\mathcal{S}_1| + 1$ times, every edge incident to b is subdivided at least $2^{\Delta-1}(|\mathcal{S}_2| + 1)$ times, and there is a successful 3-search $\mathcal{S}_0$ on H_0 aligned to (a, b). By Lemma 59, there exists a 3-search $\mathcal{S}_a$ on H_0 aligned to (a, b) with $a \in (S_a)_t$ for each t such that $B_{H_0}(a, |\mathcal{S}_1|) \subseteq \mathsf{FC}_*(\mathcal{S}_a)$. Also, by Lemma 60, there is a 3-search $\mathcal{S}_b$ on H_0 aligned to (a, b) with $b \in (S_b)_t$ for each t such that $\mathsf{FC}_*(\mathcal{S}_b, V(H_0) \setminus B_{H_0}(b, |\mathcal{S}_2| + 1)) = V(H_0)$.

Let H_3 be a path from c to d of length $|\mathcal{S}_0| + 5$ that is disjoint from $H_0 \cup H_1 \cup H_2$ except for the vertices c and d. The desired subdivision of G is the union $H_0 \cup H_1 \cup H_2 \cup H_3$.

Let $\mathcal{S}_{pa}$ be a 3-search on $H_0 \cup H_1 \cup H_2 \cup H_3$ of length $|\mathcal{S}_0| + 2$, where $(S_{pa})_t$ comprises the vertices of distance $t - 1$ and t from c in H_3, as well as a. Similarly, let $\mathcal{S}_{pb}$ be a search of length $|\mathcal{S}_0| + 2$, where $(S_{pb})_t$ comprises the vertices of distance $t + 2$ and $t + 3$ from c in H_3, as well as b. (In particular, the last set in $\mathcal{S}_{pb}$ contains d.) The desired successful 3-search on $H_0 \cup H_1 \cup H_2 \cup H_3$ is the concatenation

$$\mathcal{S}_a{}^\frown\mathcal{S}_1{}^\frown\mathcal{S}_{pa}{}^\frown\mathcal{S}_0{}^\frown\mathcal{S}_{pb}{}^\frown(\{d\}){}^\frown\mathcal{S}_2{}^\frown\mathcal{S}_b.$$

Notice that in this search, after performing $\mathcal{S}_{pb}$ we just check the one-element set $\{d\}$. This is not strictly speaking necessary but makes the analysis a bit simpler.

Let H be the disjoint union of H_0, H_1, H_2, and H_3 and denote the copy of each vertex $x \in \{b, c, d\}$ in $H_i \subset H$ by x_i. Let $\mathcal{E}$ be the equivalence relation on $V(H)$ relating a_0 with a_1, b_0 with b_2, c_1 with c_3, and d_2 with d_3. Then the quotient graph $H/\mathcal{E}$ can be naturally identified with the (non-disjoint) union of H_0, H_1, H_2, and H_3.

We construct searches $\mathcal{S}_a'$, $\mathcal{S}_b'$, $\mathcal{S}_0'$, $\mathcal{S}_1'$, $\mathcal{S}_2'$, $\mathcal{S}_{pa}'$, and $\mathcal{S}_{pb}'$ on H from $\mathcal{S}_a$, $\mathcal{S}_b$, $\mathcal{S}_0$, $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{S}_{pa}$, and $\mathcal{S}_{pb}$ respectively by replacing every appearance of b, c, or d with both of their copies in H. Define

$$\mathcal{S} = \mathcal{S}_a' \frown \mathcal{S}_1' \frown \mathcal{S}_{pa}' \frown \mathcal{S}_0' \frown \mathcal{S}_{pb}',$$

and $\mathcal{R} = (\{d_2, d_3\}) \frown \mathcal{S}_2' \frown \mathcal{S}_b'$. For brevity, define

$$t_0 = |\mathcal{S}_a|, \quad t_1 = t_0 + |\mathcal{S}_1|, \quad t_2 = t_1 + |\mathcal{S}_{pa}|, \quad t_3 = t_2 + |\mathcal{S}_0|, \quad t_4 = t_3 + |\mathcal{S}_{pb}|, \quad t_5 = 1 + |\mathcal{S}_2|.$$

We now check that $\mathcal{S}$ is $\mathcal{E}$-invariant. For $t \leqslant t_0$, since $a \in (S_a)_t$, both a_0 and a_1 belong to S_t, and thus $a_0, a_1 \in \mathsf{PC}_t(\mathcal{S})$. It is also clear that none of b_0, b_2, c_1, c_3, d_2, d_3 are in $\mathsf{PC}_t(\mathcal{S})$ for $t \leqslant t_0$.

For $t_0 < t \leqslant t_1$, since $B_H(a_0, |\mathcal{S}_1|) \subseteq \mathsf{FC}_{t_0}(\mathcal{S})$, we have $N_H(a_0) \cup \{a_0\} \subseteq \mathsf{FC}_{t-1}(\mathcal{S})$, and so $a_0 \in \mathsf{PC}_t(\mathcal{S})$. Since $\mathcal{S}_1$ is aligned to (a, c), we also have $a_1 \in \mathsf{PC}_t(\mathcal{S})$ and $c \notin \mathsf{FC}_{t-1}(\mathcal{S})$ so

$$c_1 \in \mathsf{PC}_t(\mathcal{S}) \iff c_1 \in S_t \iff c_3 \in S_t \iff c_3 \in \mathsf{PC}_t(\mathcal{S}).$$

It is also clear that none of b_0, b_2, d_2, d_3 are in $\mathsf{PC}_t(\mathcal{S})$ for $t_0 < t \leqslant t_1$.

Since $\mathcal{S}_1$ is successful, $V(H_1) \subseteq \mathsf{FC}_{t_1}(\mathcal{S})$, and hence $\{a_1, c_1\} \subseteq \mathsf{PC}_t(\mathcal{S})$ for all $t \geqslant t_1$. Furthermore, if $t_1 < t \leqslant t_2$, then $a_0 \in \mathsf{PC}_t(\mathcal{S})$ since $\{a_0, a_1\} \subseteq S_t$. Note also that if $t_1 < t \leqslant t_2$, then $\mathsf{PC}_t(\mathcal{S}) \cap V(H_3) = B_H(c_3, t - t_1)$, so $c_3 \in \mathsf{PC}_t(\mathcal{S})$. Finally, none of b_0, b_2, d_2, d_3 are in $\mathsf{PC}_t(\mathcal{S})$ for $t_1 < t \leqslant t_2$.

For $t_2 < t \leqslant t_3$, we have $\{a_1, c_1\} \subseteq \mathsf{PC}_t(\mathcal{S})$. Since $\mathcal{S}_0$ is aligned to (a, c), $a_0 \in \mathsf{PC}_t(\mathcal{S})$ and

$$b_0 \in \mathsf{PC}_t(\mathcal{S}) \iff b_0 \in S_t \iff b_2 \in S_t \iff b_2 \in \mathsf{PC}_t(\mathcal{S}).$$

Since $B_H(c_3, |\mathcal{S}_0| + 1) \subseteq \mathsf{FC}_{t_2}(\mathcal{S})$, we have $c_3 \in \mathsf{PC}_t(\mathcal{S})$. Finally, neither d_2 nor d_3 are in $\mathsf{PC}_t(\mathcal{S})$.

For $t_3 < t \leqslant t_4$, we have $\{a_1, c_1\} \subseteq \mathsf{PC}_t(\mathcal{S})$. Since $\mathcal{S}_0$ is successful, $V(H_0) \subseteq \mathsf{PC}_{t_3}(\mathcal{S})$, and hence $\{a_0, b_0\} \subseteq \mathsf{PC}_t(\mathcal{S})$. We note that $\mathsf{FC}_{t_2}(\mathcal{S})$ contains $B_H(c_3, |\mathcal{S}_0| + 1)$ by construction, so $\mathsf{FC}_{t_3}(\mathcal{S})$ contains $B_H(c_3, 1)$. We then see that $\mathsf{FC}_t(\mathcal{S})$ includes $B_H(c_3, 1 + t - t_3)$, and in particular it contains c_3. Both b_0 and b_2 are in $\mathsf{PC}_t(\mathcal{S})$ since $\{b_0, b_1\} \subseteq S_t$. Finally, it is clear that

$$d_2 \in \mathsf{PC}_t(\mathcal{S}) \iff d_3 \in \mathsf{PC}_t(\mathcal{S}) \iff t = t_4.$$

We have thus shown the $\mathcal{E}$-invariance of $\mathcal{S}$. It is straightforward to check that $\mathsf{FC}_*(\mathcal{S}) = V(H_0) \sqcup V(H_1) \sqcup V(H_3)$, so, by Proposition 58, $\vee_\mathcal{E}(\mathcal{S})$ is a 3-search on $H/\mathcal{E}$ with

$$\mathsf{FC}_*(\vee_\mathcal{E}(\mathcal{S})) = V(H_0) \cup V(H_1) \cup V(H_3) \setminus \{b, d\}$$

(here the union is non-disjoint). Let

$$A = V(H_0) \sqcup V(H_1) \sqcup V(H_3) \setminus \{b_0, d_3\}.$$

Then A is an $\mathcal{E}$-invariant subset of $V(H)$ and $A/\mathcal{E} = \mathsf{FC}_*(\vee_\mathcal{E}(\mathcal{S}))$. We now wish to show that the search $\mathcal{R}$ is $\mathcal{E}$-invariant over A.

Note that $\{a_0, a_1, c_1, c_3, d_2, d_3\} \subseteq \mathsf{PC}_1(\mathcal{R}, A)$ and $\{b_0, b_2\} \cap \mathsf{PC}_1(\mathcal{R}, A) = \varnothing$. Furthermore, $\mathsf{FC}_1(\mathcal{R}, A) = V(H_0) \sqcup V(H_1) \sqcup V(H_3) \setminus B_H(b_0, 1)$.

For $1 < t \leqslant t_5$, $\{a_0, a_1, c_1, c_3, d_3\} \subseteq \mathsf{FC}_t(\mathcal{R}, A)$. Since $\mathcal{S}_2$ is aligned to (d, b), $d_2 \in \mathsf{PC}_t(\mathcal{R}, A)$ and

$$b_0 \in \mathsf{PC}_t(\mathcal{R}, A) \iff b_0 \in R_t \iff b_2 \in R_t \iff b_2 \in \mathsf{PC}_t(\mathcal{R}, A).$$

Since $\mathcal{S}_2$ is successful, $V(H_2) \subseteq \mathsf{FC}_{t_5}(\mathcal{R}, A)$.

For $t > t_5$, we have $\{a_1, c_1, b_2, d_2, c_3, d_3\} \subseteq \mathsf{FC}_t(\mathcal{R}, A)$. Since $\mathcal{S}_b$ is aligned to (a, b), we have $a_0 \in \mathsf{PC}_t(\mathcal{R}, A)$, and by construction $b_0 \in R_t \subseteq \mathsf{PC}_t(\mathcal{R}, A)$.

We have thus shown that $\mathcal{R}$ is $\mathcal{E}$-invariant over A. Hence, $\vee_\mathcal{E}(\mathcal{R})$ is a successful 3-search with $A/\mathcal{E}$ as its initial cleared set. It is also clear that $\vee_\mathcal{E}(\mathcal{R})$ is aligned to (a, b), so $\vee_\mathcal{E}(\mathcal{S}) \frown \vee_\mathcal{E}(\mathcal{R})$ is a successful 3-search on $H/\mathcal{E}$ aligned to (a, b) as desired. ∎

9.5 Searching graphs with simple GSP decompositions

We are now ready to complete the proof of the implication $(3) \implies (1)$ in Theorem 30.

Theorem 65. *Suppose G is a connected graph with a simple GSP decomposition. Then $\mathsf{in}_t(G) \leqslant 3$.*

Proof. We shall show by induction on $|V(G)|$ that if $\boldsymbol{G} = (G, a, b)$ admits a simple GSP decomposition T, then every subdivision of G has a further subdivision which has a successful 3-search aligned to (a, b). The base case, where G is a single edge ab, is clear.

Now assume the GSP decomposition T is nontrivial. If $\circ_{T,\boldsymbol{G}} \in \{\circ_s, \circ_b\}$, then we are done by Propositions 61 and 62 respectively and the inductive hypothesis. If $\circ_{T,\boldsymbol{G}} = \circ_{b'}$, i.e., $\boldsymbol{G} = \boldsymbol{G_0} \circ_{b'} \boldsymbol{G_1}$, then we can use Lemma 47 to switch the order of the terminals in $\boldsymbol{G_1}$ and then apply Proposition 63 and the inductive hypothesis.

It remains to consider the case when $\circ_{T,\boldsymbol{G}} = \circ_p$, i.e., $(G, a, b) = (G_0, a, b) \circ_p (K, a, b)$. Since T is simple, we may assume, without loss of generality, that K has a bridge cd separating a and b, with c closer to a and d closer to b. To apply Proposition 64, we need to decompose (K, a, b) as $(G_1, a, c) \circ_s (e_{cd}, c, d) \circ_s (G_2, d, b)$, where e_{cd} is the graph comprising the single edge cd.

Claim 65.a. *Suppose (K, a, b) contains a bridge cd separating its terminals (with c closer to a and d closer to b) and admits a simple GSP decomposition T'. Let G_0 be the subgraph of K induced by the vertices that are separated from b by d, and let G_1 be the subgraph of K induced by the vertices that are separated from a by c. Then (G_0, a, c) and (G_1, d, b) each either comprise a single vertex or admit simple GSP decompositions.*

▷   We shall induct on $|V(K)|$. The base case where K comprises a single edge is trivial. Now suppose the decomposition T' is nontrivial. Note that $\circ_{T',\mathbf{K}} \neq \circ_p$, since otherwise K would contain two internally disjoint ab-paths. If $\circ_{T',\mathbf{K}} = \circ_b$, then $(K, a, b) = (K_0, a, b) \circ_b (K_1, a, x)$ for some vertex x. Then cd is a bridge separating the terminals of (K_0, a, b), so, by the inductive hypothesis, the subgraphs H_0 and H_1 of K_0 comprising the vertices separated from each terminal by d, c respectively have simple GSP decompositions. But $G_1 = H_1$ and $(G_0, a, c) = (H_0, a, c) \circ_b (K_1, a, x)$. This yields simple GSP decompositions of G_0 and G_1, as desired. The case where $\circ_{T',\mathbf{K}} = \circ_{b'}$ follows similarly, *mutatis mutandis*.

Finally, suppose $\circ_{T',\mathbf{K}} = \circ_s$, i.e., $(K, a, b) = (K_0, a, x) \circ_s (K_1, x, b)$ for some vertex x. Since c and d are adjacent, they must both be contained in either K_0 or K_1; without loss of generality, say $c, d \in V(K_0)$. By the inductive hypothesis, the subgraphs H_0 and H_1 of K_0 comprising the vertices separated from each terminal by d, c respectively have simple GSP decompositions. But $G_0 = H_0$, and $(G_1, d, b) = (H_1, d, x) \circ_s (K_1, x, b)$, which yields simple GSP decompositions of G_0 and G_1, as desired. ◀

If necessary, we may replace K with a subdivision and assume that $\{a, b\} \cap \{c, d\} = \varnothing$. By the above claim, we may then write $(K, a, b) = (G_0, a, c) \circ_s (e_{cd}, c, d) \circ_s (G_1, d, b)$, where (G_0, a, c) and (G_1, d, b) admit simple GSP decompositions. We can then apply Proposition 64 to finish the proof. ∎

10   Further investigations

Our main result, Theorem 7, provides a complete characterization of graphs with topological inspection number at most 3. A natural next step would be to consider graphs with topological inspection number at most k for larger values of k. In this case, it is perhaps too optimistic to seek a complete characterization; however, even many rather fundamental specific questions remain open. For instance, we do not know the values of in_t for such basic families of graphs as complete graphs or complete bipartite graphs:

Open Problem 66. Determine $\mathsf{in}_t(K_n)$ and $\mathsf{in}_t(K_{s,t})$.

Since $\mathsf{in}_t(K_2) = 2$, $\mathsf{in}_t(K_3) = 3$, and (by Theorem 32) $\mathsf{in}_t(K_4) = 4$, a natural guess would be that $\mathsf{in}_t(K_n) = n$ for all n. This, however, turns out not to be the case, as the following bound shows:

Proposition 67. $\mathsf{in}_t(K_n) \leqslant \lceil \frac{n}{3} \rceil + 2$.

Proof sketch. Given any subdivision G of K_n, partition the vertices of K_n in G into three parts A, B, C, each of size at most $\lceil \frac{n}{3} \rceil$. Let $\mathcal{E}$ be the equivalence relation on $V(G)$ where A, B, and C are the nontrivial equivalence classes (that is, every vertex of G that is not one of the original vertices of K_n forms its own equivalence class). Now consider $G/\mathcal{E}$. This graph comprises three vertices a, b, c, which are the images of A, B, C respectively under the quotient map, along with several internally vertex-disjoint paths between each of them, as well as some cycles which only intersect the rest of the graph at one of a, b, or c. Let P_{ab}, P_{bc}, P_{ac} be the subgraphs of $G/\mathcal{E}$ comprising the paths between their respective

subscripted vertices, and let C_a, C_b, C_c be the subgraphs comprising the cycles containing their respective subscripted vertices.

We will construct a successful 3-search on some subdivision of $G/\mathcal{E}$ in which at most one of a, b, or c is contained in each search set. Note that

$$(G/\mathcal{E}, a, c) = (P_{ac}, a, c) \circ_p ((P_{ab}, a, b) \circ_s ((P_{bc}, b, c) \circ_b (C_b, b, y)) \circ_b (C_a, a, x) \circ_{b'} (C_c, c, z),$$

where x, y, and z are arbitrary vertices of C_a, C_b, and C_c respectively distinct from a, b, and c. By Theorem 65, every subdivision of each of the graphs (P_{ac}, a, c), (P_{ab}, a, b), (P_{bc}, b, c), (C_b, b, y), (C_a, a, x), (C_c, c, z) admits a further subdivision that has a successful 3-search aligned to the corresponding terminals; furthermore, it is easy to verify that if the graphs are sufficiently subdivided, then the searches constructed in the proof of Theorem 65 include at most one terminal vertex in each search set. Combining these searches via Propositions 61, 62, 63, and 64 gives a successful 3-search $\mathcal{S}$ on some subdivision of $G/\mathcal{E}$ in which at most one of a, b, or c is contained in each search set, as desired.

Now note that every subdivision of $G/\mathcal{E}$ is the quotient graph under $\mathcal{E}$ of some subdivision of G. Let G' be the subdivision of G such that $G'/\mathcal{E}$ is the graph on which $\mathcal{S}$ is defined. Let $\mathcal{S}'$ be the search on G' obtained by taking $\mathcal{S}$ and replacing each instance of a, b, or c with the corresponding equivalence class in $\mathcal{E}$. Since at most one of a, b, or c is in each S_t, we have $|S_t'| \leqslant \lceil \frac{n}{3} \rceil + 2$ and thus $\mathcal{S}'$ is a $(\lceil \frac{n}{3} \rceil + 2)$-search. By construction, $\mathcal{S}'$ is $\mathcal{E}$-invariant, and thus by Proposition 58 it is successful, proving the claim. ∎

In particular, $\mathsf{in}_t(K_5) = \mathsf{in}_t(K_6) = 4$. It also follows that $4 \leqslant \mathsf{in}_t(K_7) \leqslant 5$, but we do not know which one is the correct value. We conjecture that $\lim_{n \to \infty} \mathsf{in}_t(K_n) = \infty$, but the best lower bound that we can currently prove is 4. Indeed, we do not know is there exist *any* graphs G with $\mathsf{in}_t(G) > 4$!

Open Problem 68. Is there $n \in \mathbb{N}$ such that $\mathsf{in}_t(K_n) > 4$? Is it true that in fact $\lim_{n \to \infty} \mathsf{in}_t(K_n) = \infty$?

Another interesting class of graphs to consider comprises planar graphs. It follows from Theorem 7 that outerplanar graphs have topological inspection number at most 3, and that some planar graphs (e.g., K_4) have topological inspection number 4. We do not know if there is a constant upper bound that works for all planar graphs:

Open Problem 69. Is there a constant k such that $\mathsf{in}_t(G) \leqslant k$ for all planar graphs G? Does $k = 4$ work?

We finish the paper with an enticing conjecture that may provide a characterization of the inspection number in terms of vertex sets with small boundary. Observe that every lower bound on the inspection number established in this paper is proved using Proposition 12 (this includes Theorems 2, 15, 32, 38, and 39). Namely, to show that $\mathsf{in}(G) > k$, we exhibit a subgraph H of G that does not contain a subset $C \subseteq V(H)$ of size close to a certain value and with $|\partial_H(C)| < k$. Since this is the only method we have for proving lower bounds on $\mathsf{in}(G)$, it is natural to wonder if Proposition 12 has a converse:

Open Problem 70. If $\mathsf{in}(G) > k$, do there necessarily exist a subgraph $H \subseteq G$ and an integer $1 \leqslant i \leqslant |V(H)|$ such that there is no subset $C \subseteq V(H)$ with $|\partial_H(C)| < k$ and $i - k < |C| < i$?

References

[1] T. Abramovskaya, F. Fomin, P. Golovach, and M. Pilipczuk. *How to hunt an invisible rabbit on a graph*, Eur. J. Comb., **52** (A) (2016), 12–26 (cit. on p. 2)

[2] B. Alspach. *Searching and sweeping graphs: a brief survey*, Le Matematiche, **59** (1,2) (2004), 5–37 (cit. on p. 2)

[3] H. Aydinian, F. Cicalese, C. Deppe, and V. Lebedev. *A combinatorial model of two-sided search*, Int. J. Found. Comput. Sci., **29** (04) (2018), 481–504 (cit. on p. 2)

[4] F. Berger, A. Gilbers, A. Grüne, and R. Klein. *How many lions are needed to clear a grid?*, Algorithms, **2** (3) (2009), 1069–1086 (cit. on p. 2)

[5] D. Bienstock and P. Seymour. *Monotonicity in graph searching*, Journal of Algorithms, **12** (2) (1991), 239–245 (cit. on p. 2)

[6] A. Bonato and B. Yang. *Graph searching and related problems*, Handbook of Combinatorial Optimization, (2013). Ed. by P. Pardalos, D.Z. Du, and R. Graham, 1511–1558 (cit. on p. 2)

[7] J.R. Britnell and M. Wildon. *Finding a princess in a palace: A pursuit-evasion problem*, Electron. J. Comb., **20** (1) (2013) (cit. on p. 2)

[8] J. Carraher, I. Choi, M. Delcourt, L.H. Erickson, and D.B. West. *Locating a robber on a graph via distance queries*, Theor. Comput. Sci., **463** (2012), 54–61 (cit. on p. 2)

[9] D. Dereniowski, D. Dyer, R.M. Tifenbach, and B. Yang. *The complexity of zero-visibility cops and robber*, Theor. Comput. Sci., **607** (2) (2015), 135–148 (cit. on p. 2)

[10] R. Diestel. *Graph Theory*. 5th ed. Berlin Heidelberg: Springer-Verlag, 2017 (cit. on pp. 19, 22)

[11] R.J. Duffin. *Topology of series-parallel networks*, J. Math. Anal. Appl., **10** (2) (1965), 303–318 (cit. on p. 4)

[12] J. Ellis and R. Warren. *Lower bounds on the pathwidth of some grid-like graphs*, Discret. Appl. Math., **156** (5) (2008), 545–555 (cit. on p. 2)

[13] F.V. Fomin and D.M. Thilikos. *An annotated bibliography on guaranteed graph searching*, Theor. Comput. Sci., **399** (3) (2008), 236–245 (cit. on p. 2)

[14] J. Haslegrave. *An evasion game on a graph*, Disc. Math., **314** (2014), 1–5 (cit. on p. 2)

[15] A.S. LaPaugh. *Recontamination does not help to search a graph*, J. ACM, **40** (2) (1993), 224–245 (cit. on p. 2)

[16] R. Tošić. *Vertex-to-vertex search in a graph*, Graph theory (Dubrovnik, 1985), (1986), 233–237 (cit. on p. 2)

[17] B. Yang, D. Dyer, and B. Alspach. *Sweeping graphs with large clique number*, Discrete Math., **309** (18) (2009), 5770–5780 (cit. on p. 2)