

On an Induced Version of Menger's Theorem

Kevin Hendrey^a Sergey Norin^b Raphael Steiner^c
J eremie Turcotte^b

Submitted: Nov 23, 2023; Accepted: Sep 11, 2024; Published: Nov 1, 2024

  The authors. Released under the CC BY-ND license (International 4.0).

Abstract

We prove Menger-type results in which the obtained paths are pairwise non-adjacent, both for graphs of bounded maximum degree and, more generally, for graphs excluding a topological minor. More precisely, we show the existence of a constant C , depending only on the maximum degree or on the forbidden topological minor, such that for any pair of sets of vertices X, Y and any positive integer k , there exist either k pairwise non-adjacent X - Y -paths, or a set of fewer than Ck vertices which separates X and Y . We further show better bounds in the subcubic case, and in particular obtain a tight result for two paths using a computer-assisted proof.

Mathematics Subject Classifications: Primary 05C38; Secondary 05C15, 05C40, 05C83

1 Introduction

Given a graph G and $X, Y \subseteq V(G)$, we say a set of vertices Z *separates* X and Y if Z intersects every X - Y -path¹. In general, we say two paths are disjoint if they do not share any vertices. Menger's theorem is a fundamental result of graph theory, relating the existence of many disjoint paths between two sets of vertices in a graph with the absence of small separators.

Theorem 1 (Menger's theorem [12]). *If $k \in \mathbb{N}$, G is a graph and $X, Y \subseteq V(G)$, then there exist either*

^aDiscrete Mathematics Group, Institute for Basic Science (IBS), Daejeon, South Korea
(kevinhendrey@ibs.re.kr).

^bDepartment of Mathematics and Statistics, McGill University, Montr al, Canada
(snorin@math.mcgill.ca, mail@jeremieturcotte.com).

^cInstitute of Theoretical Computer Science, Department of Computer Science, ETH Z rich, Switzerland (raphaelmario.steiner@inf.ethz.ch).

¹An X - Y -path is defined as a path that has one endpoint in X and the opposite endpoint in Y . Notably, this definition also allows a single vertex in $X \cap Y$ to qualify as an X - Y -path (of length 0).

- (1) k pairwise disjoint X - Y -paths, or
- (2) a set of fewer than k vertices which separates X and Y .

It is a natural question to ask under which circumstances we can guarantee the paths in point (1) to be far apart from each other in the graph metric, rather than just disjoint. Georgakopoulos and Papasoglu [6], motivated by questions in metric geometry, and Albrechtsen et al. [1] conjectured a “Coarse Menger’s theorem”. Although it was later disproved by Nguyen et al. [14], we will state it formally, as it motivates our work and various relaxations of it are still open. Before stating it, we need the following notation. If G is a graph, $Z \subseteq V(G)$ and $d \in \mathbb{N}$, we write $B_G(Z, d)$ for the *ball of radius d around Z* , i.e. all vertices at distance at most d from one of the vertices in Z . If $Z = \{z\}$, we may simply write $B_G(z, d)$.

Conjecture 2. For every $k \in \mathbb{N}$, there exists $c = c(k) \in \mathbb{N}$ satisfying the following. If $d \in \mathbb{N}$, G is a graph and $X, Y \subseteq V(G)$, then there exist either

- (1) k disjoint X - Y -paths P_1, \dots, P_k such that $\text{dist}_G(P_i, P_j) \geq d$ for all distinct i, j , or
- (2) a set Z of fewer than k vertices such that $B_G(Z, cd)$ separates X and Y .

We note that the version of the conjecture of Albrechtsen et al. was in fact stronger, as it did not allow c to depend on k . It is also interesting to note that McCarty and Seymour had shown that it was sufficient to prove the conjecture for the case $d = 3$ for the entire conjecture to hold, see [1, Theorem 4].

Both Georgakopoulos and Papasoglu [6] and Albrechtsen et al. [1] have shown that conjecture 2 holds for $k = 2$, the constant of 129 below is from the latter authors.

Theorem 3 ([6, 1]). *If $d \in \mathbb{N}$, G is a graph and $X, Y \subseteq V(G)$, then there exist either*

- (1) two disjoint X - Y -paths P_1, P_2 such that $\text{dist}_G(P_1, P_2) \geq d$, or
- (2) $z \in V(G)$ such that $B_G(z, 129d)$ separates X and Y .

One possible variation on conjecture 2 is to replace the family of balls $B_G(Z, cd)$ which separates X and Y in condition (2) by an arbitrary set. What is a natural upper bound on the size of this set? If G has maximum degree bounded by a constant Δ , one can upper-bound the size of $B_G(Z, cd)$ by $|B_G(Z, cd)| \leq |Z| \cdot \sum_{i=0}^{cd} \Delta^i < \Delta^{cd+1}k$. In particular, if the strong version of conjecture 2 proposed by Albrechtsen et al. had held (that is, with c independent of k), it would have implied the following.

Conjecture 4. For every $d, \Delta \in \mathbb{N}$ there exists a constant $C = C(d, \Delta) > 0$ such that the following holds. If $k \in \mathbb{N}$, G is a graph with $\Delta(G) \leq \Delta$ and $X, Y \subseteq V(G)$, then there exist either

- (1) k X - Y -paths pairwise at distance at least d in G , or
- (2) a set of fewer than Ck vertices in G which separates X and Y .

We note that the reduction for $d \geq 3$ to $d = 3$ by McCarty and Seymour also holds for conjecture 4, see appendix A.

In the first main result of this paper, namely theorem 5 below, we prove conjecture 4 in the first non-trivial case when $d = 2$, that is, when we look for a family of disjoint X - Y -paths that are pairwise non-adjacent². For brevity, when saying that paths are pairwise non-adjacent, they will also be meant to be disjoint.

To state theorem 5 concisely we need a bit of terminology. Let G be a graph. We say $M \subseteq E(G)$ is an *induced matching* if $\text{dist}_G(e_1, e_2) \geq 2$ for every distinct $e_1, e_2 \in M$. A *strong edge colouring* of G is a partition of the edges of G into induced matchings. In other words, the edges are coloured such that no two edges of the same colour are adjacent. The *strong chromatic index* of G , denoted by $\chi'_s(G)$, is the smallest number of matchings in a strong edge colouring of G . The strong chromatic index is well studied, and there are many known bounds depending on the maximum degree $\Delta(G)$ of G , as well as for more specific classes of graphs. In general, $\chi'_s(G) \leq 2\Delta(G)(\Delta(G) - 1) + 1$, which can be seen by counting the number of edges at distance at most 2 of any edge and colouring greedily. For large enough Δ , currently the best bound, by Hurley, de Verclos and Kang [10], is $\chi'_s(G) \leq 1.772\Delta^2$ when $\Delta(G) \leq \Delta$.

Theorem 5. *If $k \in \mathbb{N}$, G is a graph and $X, Y \subseteq V(G)$, then there exist either*

- (1) *k pairwise non-adjacent X - Y -paths, or*
- (2) *a set of fewer than $2^{\chi'_s(G)}k$ vertices which separates X and Y .*

Given the bounds on $\chi'_s(G)$ mentioned above, we may obtain the $d = 2$ case of conjecture 4 as a direct corollary of theorem 5.

By Menger's theorem (theorem 1), theorem 5 is equivalent to the following result.

Theorem 6. *If $k \in \mathbb{N}$, G is a graph, $X, Y \subseteq V(G)$ and there exist at least $2^{\chi'_s(G)}k$ pairwise disjoint X - Y -paths, then there exist k pairwise non-adjacent X - Y -paths in G .*

We believe this result to be interesting in its own right, as it is a quite natural analogue to Menger's theorem in an induced setting. An example of another result of this type would be Korhonen's [11] proof of the grid minor theorem for induced minors for bounded degree graphs.

The idea behind our proof is as follows. Given a large number of disjoint X - Y -paths and a strong colouring of the edges in between the paths, we contract all edges of a colour class (say, green) and apply Menger's theorem to find many disjoint X - Y -paths in the contracted graph, which we can then lift back to the original graph. By this contraction, we will be guaranteed to not have any green edges between the paths, and the strong colouring will guarantee that there are no edges of the original paths that go between the

²Since for every X - Y -path P in a graph G there exists an induced X - Y -path P' such that $V(P') \subseteq V(P)$, one can see that the existence of a family of k pairwise non-adjacent X - Y -paths is equivalent to the existence of a family \mathcal{P} of k different X - Y -paths, such that the union of the paths in \mathcal{P} forms an induced subgraph of G . This explains the naming of the paper.

new paths. After repeating this argument for every colour, we find a collection of pairwise non-adjacent paths.

We will prove theorem 6 in section 2. In fact, given that we do not need to colour the edges of the original paths, we can obtain an improvement over the constant $2^{\chi'_s(G)}$, which is most significant when the maximum degree is small. In particular, in section 4, we will show the following two results. For brevity, if \mathcal{P} is a collection of subgraphs (typically, of paths) of a graph G , then write $V(\mathcal{P}) = \bigcup_{P \in \mathcal{P}} V(P)$ and $E(\mathcal{P}) = \bigcup_{P \in \mathcal{P}} E(P)$.

Theorem 7. *If G is a graph, $X, Y \subseteq V(G)$, and there exists a collection \mathcal{P} of at least $16k$ disjoint X - Y -paths such that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$, then there exist at least k pairwise non-adjacent X - Y -paths in G .*

This is a large improvement over the constant 2^{10} which would be obtained from theorem 6 with the bound of $\chi'_s(G) \leq 10$ when $\Delta(G) \leq 3$ proved independently by Andersen [2] and by Horák, Qing and Trotter [9].

For the case $k = 2$, we will show in a computer-assisted proof the following tight result.

Theorem 8. *If G is a graph, $X, Y \subseteq V(G)$, and there exists a collection \mathcal{P} of five disjoint X - Y -paths such that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$, then there exist two pairwise non-adjacent X - Y -paths in G .*

Furthermore, the statement does not necessarily hold if either

- (a) \mathcal{P} contains four paths instead of five, or
- (b) we replace the condition that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$ by the condition that the maximum degree of G is three.

Both the proof methods of Theorems 7 and 8 and (b) in the latter indicate that the maximum degree of $G - E(\mathcal{P})$ is a more natural parameter to bound than the maximum degree of G .

A direct consequence of (a) in this result is that the $16k$ in theorem 7 cannot be improved below $4k - 3$ (consider a disjoint union of $k - 1$ copies of a graph containing 4 disjoint X - Y -paths but no two non-adjacent X - Y -paths).

We say a graph H is a *topological minor* of a graph G if G contains a subdivision³ of H as a subgraph. In section 3, using the structure theorem for graphs excluding a topological K_r -minor first proved by Grohe and Marx [7], as well as Erde and Weißbauer [4], we then generalize our induced Menger's theorem to the class of graphs excluding the complete graph K_r as a topological minor.

Theorem 9. *For every $r \in \mathbb{N}$, there exists $c = c_9(r) > 0$ such that the following holds.*

If G is a graph not containing K_r as a topological minor and $X, Y \subseteq V(G)$, then there exist either

- (1) k pairwise non-adjacent X - Y -paths, or

³As usual, by a *subdivision* of H we here mean any graph that is isomorphic to a graph that can be obtained from H by replacing its edges by paths of positive length.

(2) a set of fewer than ck vertices which separates X and Y .

We note that theorems 5 and 9 were also proved, with slightly different constants, independently by Gartland et al. [5].

We will need the following notation. If G is a graph and S is a set of vertices of G , then $G[S]$ is the graph subgraph of G induced on S , and $G - S$ is the subgraph of G induced on $V(G) \setminus S$. If M is a set of edges of G , then $G - M$ is the subgraph of G with vertex set $V(G)$ and set of edges $E(G) \setminus M$.

2 Graphs with bounded maximum degree

theorem 6 follows directly from the following result, by taking \mathcal{M} as the colour classes of edges in $E(G[V(\mathcal{P})]) \setminus E(\mathcal{P})$ (in other words, the edges not in \mathcal{P} but with both ends in vertices in \mathcal{P}) in a strong edge colouring of G .

Theorem 10. *If $m, k \in \mathbb{N}$, G is a graph, $X, Y \subseteq V(G)$, \mathcal{P} is a collection of $2^m k$ pairwise disjoint X - Y -paths and \mathcal{M} is a partition of $E(G[V(\mathcal{P})]) \setminus E(\mathcal{P})$ into at most m induced matchings of G , then there exist k pairwise non-adjacent X - Y -paths in G .*

Proof. First note that we may without loss of generality assume that $V(G) = V(\mathcal{P})$, by restricting G, X, Y to $V(\mathcal{P})$. This implies that $E(G) = E(\mathcal{P}) \cup (\cup \mathcal{M})$, where $\cup \mathcal{M} := \cup_{M \in \mathcal{M}} M$.

We prove the statement by induction on m . If $m = 0$, then $\mathcal{M} = \emptyset$. In particular, $E(G[V(\mathcal{P})]) \setminus E(\mathcal{P}) = \emptyset$, and so the $2^0 k = k$ paths of \mathcal{P} are pairwise non-adjacent.

We now show the inductive step. We may assume that the paths of \mathcal{P} are chosen such that the sum of the lengths of paths in \mathcal{P} is smallest possible among all collections of k disjoint X - Y -paths. This immediately implies that every $P \in \mathcal{P}$ is an induced path in G , and that every path $P \in \mathcal{P}$ intersects X and Y only in its endpoints.

Let $M \in \mathcal{M}$, chosen arbitrarily, and write $\mathcal{M}^* := \mathcal{M} \setminus \{M\}$. We define $G' = G - E(\mathcal{M}^*)$, and let G'' be obtained from G' by contracting the edges of M .⁴ Let $f : V(G) \rightarrow V(G'')$ denote the mapping of vertices underlying the resulting contraction. The sets corresponding to X, Y in G'' are then $X' := f(X), Y' := f(Y)$.

By Menger's theorem (theorem 1), there exists either

1. a collection \mathcal{P}' of $2^{m-1}k$ disjoint X' - Y' -paths in G'' , or
2. a set Z' of fewer than $2^{m-1}k$ vertices separating X' and Y' in G'' .

First suppose we are in case (2). Let $Z := f^{-1}(Z')$. As we have contracted by a matching, the preimage of any vertex in G'' is of size at most two, and so $|Z| \leq 2|Z'| < 2^m k$. We claim Z separates X and Y in G' , which would be a contradiction as there are at least $2^m k$ pairwise disjoint X - Y -paths in G' (the paths in \mathcal{P} are not affected by removing M^*). If P is an X - Y -path in G' , then $f(V(P))$ corresponds to the vertex set of

⁴An edge is *contracted* by identifying its end vertices, and removing any resulting loops and parallel edges.

a $X'-Y'$ walk in G'' , from which we can extract an $X'-Y'$ -path P' . Given that Z' separates X' and Y' in G'' , there exists $v' \in Z' \cap V(P') \neq \emptyset$. By construction of P' , there exists $v \in V(P)$ such that $f(v) = v'$. By definition of Z , $v \in Z$. Hence, $v \in Z \cap V(P)$ as desired. Therefore, we are necessarily in case (1).

Given a path $P' \in \mathcal{P}'$, it is easily seen that $G'[f^{-1}(V(P'))]$ is connected and so there exists an $X-Y$ path P in G' such that $f(V(P)) \subseteq V(P')$. There may be multiple such paths, but we will choose one and refer to it as the *lift* of P' . In particular, we may suppose that this path is induced in G' . Let \mathcal{P}_2 be the collection of lifts of paths of \mathcal{P}' . Given that the paths of \mathcal{P}' are pairwise disjoint, so are those in \mathcal{P}_2 .

Let \mathcal{M}_2 be the collection of matchings of \mathcal{M}^* restricted to edges in $E(G[V(\mathcal{P}_2)]) \setminus E(\mathcal{P}_2)$.

We claim \mathcal{M}_2 partitions the edges of $E(G[V(\mathcal{P}_2)]) \setminus E(\mathcal{P}_2)$. The fact that the matchings in \mathcal{M}_2 are pairwise disjoint is direct from their construction as restrictions of matchings in \mathcal{M}^* . Let $e = uv \in E(G[V(\mathcal{P}_2)]) \setminus E(\mathcal{P}_2)$, and suppose for a contradiction that e is not in any matching of \mathcal{M}_2 .

First suppose that $e \in M$. We cannot have $u, v \in V(P)$ for $P \in \mathcal{P}_2$, since $e \notin E(\mathcal{P}_2)$, $M \subseteq E(G')$ and P is induced in G' . Hence, $u \in V(P_1)$ and $v \in V(P_2)$ for distinct paths $P_1, P_2 \in \mathcal{P}_2$. However, P_1, P_2 are lifts of paths in G'' , say $P'_1, P'_2 \in \mathcal{P}'$. In particular, $f(u) \in V(P'_1)$ and $f(v) \in V(P'_2)$. As $uv \in M$, $f(u) = f(v)$, and so P'_1 and P'_2 are not disjoint, which is a contradiction to the choice of \mathcal{P}' . Hence, we may now suppose that $e \notin M$.

Given that $e \notin M$ and e is not in any matching of \mathcal{M}^* , $e \notin \cup \mathcal{M}$. By our first assumption, necessarily $e \in E(\mathcal{P})$. If we show that both u and v are incident to some edges of M , this would be a contradiction to the fact that M is a strong matching.

We will show that u is incident to some edge of M ; the proof for v is analogous. Let $P \in \mathcal{P}_2$ be the path such that $u \in V(P)$. There are two cases to consider. First suppose u is not an endpoint of P , i.e. u has distinct neighbours $z_1, z_2 \in V(P)$. It is impossible that both $uz_1, uz_2 \in E(\mathcal{P})$, given that this is a collection of paths (so u cannot be incident to three edges of $E(\mathcal{P})$) and we know that $e \in E(\mathcal{P})$; without loss of generality say $uz_1 \in \cup \mathcal{M}$. Recall that \mathcal{P}_2 is a collection of paths in G' and so necessarily $uz_1 \in M$, as desired. Now suppose that u is an endpoint of P , hence $u \in X \cup Y$. If $u \in X \cap Y$, then e would not appear in \mathcal{P} , as we have assumed those paths to be as short as possible. Hence, $u \notin Y$ and so there exists some edge $uz \in E(\mathcal{P}_2)$. By our hypothesis that the paths in \mathcal{P} are shortest possible, u is not an interior vertex of any path in \mathcal{P} , i.e. u appears in at most one edge of \mathcal{P} , which we already know to be e . Hence $uz \notin E(\mathcal{P})$ and so $uz \in \cup \mathcal{M}$. By the same argument as previously, $uz \in M$. This completes the proof of the claim.

Note that $|\mathcal{P}_2| = |\mathcal{P}'| = 2^{m-1}k$ and $|\mathcal{M}_2| = |\mathcal{M}^*| = |\mathcal{M}| - 1 \leq m - 1$. Hence, by the induction hypothesis applied to G' , we obtain k pairwise non-adjacent $X-Y$ -paths in G , as desired. \square

We now briefly discuss why our proof does not work in the $d = 3$ case. The main difficulty is that there is no nice analogue of moving to an induced subgraph; the first step of the previous proof was that we can assume that $V(G) = V(\mathcal{P})$. Indeed, when we find non-adjacent paths in an induced subgraph of the original graph, they are also

non-adjacent in the latter. This no longer holds when $d = 3$. We would need to consider paths of length two between any paths in \mathcal{P} , but after an iteration of our method, the vertices which are on these paths of length two might now be part of \mathcal{P} . Hence, we must also consider the paths of length two between those vertices, and so on. It is unclear how one might then manage the complications arising from this.

3 Excluding a topological minor

In this section, we will prove theorem 9. We first need some definitions and notation. Let G be a graph. If $S \subseteq V(G)$, we write $G[S]$ for the subgraph of G induced by S .

A *separation* in G is to be understood as a pair (A, B) of subsets of $V(G)$ such that $A \cup B = V(G)$ and there exists no edge in G with endpoints in $A \setminus B$ and $B \setminus A$. It is slightly unusual but convenient for us to allow in this definition also degenerate cases in which $A \subseteq B$ or $B \subseteq A$. Given a separation (A, B) of G , we call $A \cap B$ its *separator* and refer to $|A \cap B|$ as the *order* of the separation (A, B) .

A *tree-decomposition* of G is a pair (T, \mathcal{V}) , where T is a tree and $\mathcal{V} = (V_t)_{t \in V(T)}$ is a collection of subsets of $V(G)$ satisfying the following properties:

- for every $v \in V(G)$, the set $\{t \in V(T) : v \in V_t\}$ induces a non-empty subtree of T , and
- for every $uv \in E(G)$, there exists at least one $t \in V(T)$ such that $u, v \in V_t$.

Given a tree decomposition (T, \mathcal{V}) of G , for every edge $e = t_1 t_2 \in E(T)$, we denote $S(e) := V_{t_1} \cap V_{t_2}$ and call $\max_{e \in E(T)} |S(e)|$ the *adhesion* of the tree-decomposition (T, \mathcal{V}) . Given a vertex $t \in V(T)$, the *torso at t* , denoted by $\tau(t)$, is defined as the graph obtained from $G[V_t]$ by adding, for every edge $e \in E(T)$ incident to t , an edge between any two non-adjacent vertices in $S(e)$, in other words we make $S(e)$ a clique for every incident edge e of t .

For every edge $e = t_1 t_2 \in E(T)$, there exists a natural corresponding separation in G , namely

$$\left(\bigcup_{t \in (T-e)(t_1)} V_t, \bigcup_{t \in (T-e)(t_2)} V_t \right),$$

where $(T - e)(t_i)$ denotes the set of vertices of the unique component of $T - e$ that contains t_i . From the definition of a tree decomposition it is not hard to see that this indeed is a separation in G , with $V_{t_1} \cap V_{t_2} = \left(\bigcup_{t \in (T-e)(t_1)} V_t \right) \cap \left(\bigcup_{t \in (T-e)(t_2)} V_t \right)$ being the corresponding separator.

Finally, we say a graph H is a *minor* of G if a graph isomorphic to H can be obtained from G by removing vertices and edges, and contracting edges. It is direct that if G contains H as a topological minor, it also contains H as a minor.

The following structure theorem is a key element of our proof of theorem 9. We use the exact statement of Erde and Weißauer [4, Theorem 4], see also Grohe and Marx [7, Theorem 4.1].

Theorem 11 ([7, 4]). *If $r \in \mathbb{N}$ and G is a graph excluding K_r as a topological minor, then G admits a tree-decomposition of adhesion less than r^2 such that every torso either*

- (1) *has fewer than r^2 vertices of degree at least $2r^4$, or*
- (2) *is K_h -minor-free, for $h = 2r^2$.*

Broadly speaking, our proof of theorem 9 will proceed as follows. Given a collection of X - Y -paths in a smallest counterexample G , we will apply theorem 11 and find a torso of the tree decomposition which intersects every path in the collection. Then, in order to find the desired collection of paths, we will either apply our result for bounded maximum degree (theorem 6), if we are in case (1), or use the following lemma, if we are in case (2).

Lemma 12. *If $h, k \in \mathbb{N}$, G is a K_h -minor-free graph, $X, Y \subseteq V(G)$ and there exist k pairwise disjoint X - Y -paths in G , then there exist at least $\frac{k}{2(h-1)}$ pairwise non-adjacent X - Y -paths in G .*

Proof. Let \mathcal{P} be a collection of k disjoint X - Y -paths in G . Let H be the minor of G obtained from $G \left[\bigcup_{P \in \mathcal{P}} V(P) \right]$ by contracting each path $P \in \mathcal{P}$ into a single vertex. In this way, the vertices of H have a natural one-to-one correspondence with the paths in \mathcal{P} , and two vertices in H are adjacent if and only if the corresponding paths in \mathcal{P} are adjacent. Since G is K_h -minor-free, so is H . Hence, by a classical result of Duchet and Meyniel [3], we have that H contains an independent set of size at least $\alpha(H) \geq \frac{|V(H)|}{2(h-1)} = \frac{k}{2(h-1)}$. The subcollection $\mathcal{P}' \subseteq \mathcal{P}$ corresponding to this independent set in H now consists of pairwise non-adjacent X - Y -paths, as desired. \square

theorem 9 follows directly from the following result, by applying Menger's theorem (theorem 1) and choosing $c_9(t) \geq \frac{1}{\varepsilon_{13}(t)}$. The additive $\frac{1}{2}$ is used solely for formal reasons, as it simplifies the inductive proof.

Theorem 13. *For every $r \in \mathbb{N}$, there exists $\varepsilon = \varepsilon(r) > 0$ such that the following holds.*

If G is a graph not containing K_r as a topological minor, $X, Y \subseteq V(G)$, $k \in \mathbb{N}$ and there are k pairwise disjoint X - Y -paths in G , then there also exists a family of at least $\varepsilon k + \frac{1}{2}$ pairwise non-adjacent X - Y -paths in G .

Proof. Fix $r \in \mathbb{N}$; we prove the statement with the constant $\varepsilon(r) := 2^{-(8r^8+3)}$. Towards a contradiction, suppose the claim is not true, and consider a counterexample G with $|V(G)|$ minimum. Hence, there exist $X, Y \subseteq V(G)$ and $k \in \mathbb{N}$ such that on the one hand, there exists a collection \mathcal{P} consisting of k pairwise disjoint X - Y -paths in G , but on the other hand, every collection \mathcal{Q} of pairwise non-adjacent X - Y -paths in G has size less than $\varepsilon k + \frac{1}{2}$. Note that the latter fact implies that $1 < \varepsilon k + \frac{1}{2}$, so $k > \frac{1}{2\varepsilon}$. Our next claim uses the minimality assumption on G to guarantee that for every separation (A, B) in G of sufficiently small order, one of its two sides must intersect all paths in \mathcal{P} .

Claim 14. *If (A, B) is a separation in G of order $|A \cap B| < 2^{8r^8+1}$, then $V(P) \cap A \neq \emptyset$ for every $P \in \mathcal{P}$ or $V(P) \cap B \neq \emptyset$ for every $P \in \mathcal{P}$.*

Proof of claim 14. Suppose towards a contradiction that there exist two paths $P_1, P_2 \in \mathcal{P}$ such that $V(P_1) \subseteq A \setminus B$ and $V(P_2) \subseteq B \setminus A$. Let $\mathcal{P}_1 := \{P \in \mathcal{P} : V(P) \subseteq A \setminus B\}$, $\mathcal{P}_2 := \{P \in \mathcal{P} : V(P) \subseteq B \setminus A\}$, and let us denote $k_1 := |\mathcal{P}_1|$, $k_2 := |\mathcal{P}_2|$. Note that since (A, B) is a separation, every path $P \in \mathcal{P} \setminus (\mathcal{P}_1 \cup \mathcal{P}_2)$ must intersect the separator $A \cap B$. Since the paths in \mathcal{P} are pairwise disjoint, this implies that $k - (k_1 + k_2) \leq |A \cap B| < 2^{8r^8+1}$. Note that \mathcal{P}_1 is a collection of $k_1 \geq 1$ pairwise disjoint $(X \cap (A \setminus B))$ - $(Y \cap (A \setminus B))$ -paths in $G[A \setminus B]$, and \mathcal{P}_2 is a collection of $k_2 \geq 1$ pairwise disjoint $(X \cap (B \setminus A))$ - $(Y \cap (B \setminus A))$ -paths in $G[B \setminus A]$. Since by our minimality assumption on G both graphs $G[A \setminus B]$ and $G[B \setminus A]$ satisfy the hypothesis of the theorem, we find that there is a collection \mathcal{Q}_1 of at least $\varepsilon k_1 + \frac{1}{2}$ pairwise non-adjacent $(X \cap (A \setminus B))$ - $(Y \cap (A \setminus B))$ -paths in $G[A \setminus B]$, and a collection \mathcal{Q}_2 of at least $\varepsilon k_2 + \frac{1}{2}$ pairwise non-adjacent $(X \cap (B \setminus A))$ - $(Y \cap (B \setminus A))$ -paths in $G[B \setminus A]$. Since there are no edges in G between $A \setminus B$ and $B \setminus A$, the collection $\mathcal{Q} := \mathcal{Q}_1 \cup \mathcal{Q}_2$ also consists of pairwise non-adjacent X - Y -paths in G . We furthermore have

$$\begin{aligned} |\mathcal{Q}| &= |\mathcal{Q}_1| + |\mathcal{Q}_2| \geq \left(\varepsilon k_1 + \frac{1}{2}\right) + \left(\varepsilon k_2 + \frac{1}{2}\right) = \varepsilon(k_1 + k_2) + 1 \\ &> \varepsilon \left(k - 2^{8r^8+1}\right) + 1 = \varepsilon k + \frac{1}{2} + \left(\frac{1}{2} - \varepsilon 2^{8r^8+1}\right) > \varepsilon k + \frac{1}{2}. \end{aligned}$$

This is a contradiction on our initial assumptions that such a collection \mathcal{Q} cannot exist. Hence our assumption was false, and this concludes the proof of the claim. \blacksquare

Next, we apply Theorem 11 to G , which yields a tree-decomposition $(T, (V_t)_{t \in V(T)})$ of G of adhesion less than r^2 , such that every torso $\tau(t)$ has at most r^2 vertices of degree at least $2r^4$, or is K_h -minor-free for $h := 2r^2$.

Claim 15. *There exists a vertex $t^* \in V(T)$ such that $V(P) \cap V_{t^*} \neq \emptyset$ for every $P \in \mathcal{P}$.*

Proof of claim 15. For every edge $e = t_1 t_2$ of T , we have that

$$\left(\bigcup_{t \in (T-e)(t_1)} V_t, \bigcup_{t \in (T-e)(t_2)} V_t \right)$$

forms a separation in G of order $|S(e)| < r^2 < 2^{8r^8+1}$. Hence, by claim 14, every path in \mathcal{P} intersects $\bigcup_{t \in (T-e)(t_1)} V_t$, or every path in \mathcal{P} intersects $\bigcup_{t \in (T-e)(t_2)} V_t$. We can therefore find an orientation \vec{T} of T such that for every edge $e = t_1 t_2$ oriented from t_1 to t_2 in \vec{T} , we have $V(P) \cap \left(\bigcup_{t \in (T-e)(t_2)} V_t\right) \neq \emptyset$ for every $P \in \mathcal{P}$. Since T is a tree, there must exist a vertex $t^* \in V(T)$ that is a sink in the orientation \vec{T} of T . We now claim that $V(P) \cap V_{t^*} \neq \emptyset$ for every $P \in \mathcal{P}$. Suppose otherwise towards a contradiction. Let $P \in \mathcal{P}$ be such that $V(P) \cap V_{t^*} = \emptyset$ and let $R := \{t \in V(T) : V(P) \cap V_t \neq \emptyset\}$. Since P is a connected subgraph of G , it readily follows from the definition of a tree-decomposition that R induces a subtree of T , which does not include t^* . Hence, there is an edge $e = t' t^* \in E(T)$ incident to t^* such

that $R \subseteq (T-e)(t')$. This, however, contradicts the fact that $V(P) \cap \left(\bigcup_{t \in (T-e)(t^*)} V_t\right) \neq \emptyset$, which follows since e is oriented from t' to t^* in \vec{T} . This concludes the proof of the claim. \blacksquare

Let H be the graph obtained from $G[V_{t^*}]$ by adding an edge between every pair x, y of non-adjacent vertices in $G[V_{t^*}]$ for which there exists a path in G with endpoints x, y all whose internal vertices are in $V(G) \setminus V_{t^*}$.

Claim 16. *For every pair of vertices $x, y \in V_{t^*}$ with $xy \notin E(G)$ for which there exists a path in G with endpoints x, y all whose internal vertices are in $V(G) \setminus V_{t^*}$, there exists an edge $f = tt^* \in E(T)$ incident with t^* such that $x, y \in S(f)$. In particular, $G[V_{t^*}] \subseteq H \subseteq \tau(t^*)$.*

Proof of claim 16. Let P be an x - y -path in G such that $V(P) \cap V_{t^*} = \{x, y\}$. Let $S := \{s \in V(T) : V_s \cap (V(P) \setminus \{x, y\}) \neq \emptyset\}$. It follows readily from the definition of a tree-decomposition (and since $P - \{x, y\}$ is a connected subgraph of G) that S induces a connected subgraph of T , i.e., $T[S]$ is a subtree of T . We furthermore have $V_{t^*} \cap (V(P) \setminus \{x, y\}) = \emptyset$, and thus $t^* \notin S$. Therefore, there exists an edge $f = tt^*$ incident with t^* such that S is contained in $(T-f)(t)$. Let x, x_1, \dots, x_ℓ, y be the vertex-trace of P . By definition of a tree-decomposition, there exist bags V_{t_1} and V_{t_2} such that $x, x_1 \in V_{t_1}$ and $x_\ell, y \in V_{t_2}$. This directly implies that $t_1, t_2 \in S \subseteq (T-f)(t)$. Hence, we have

$$x, y \in \left(\bigcup_{s \in (T-f)(t)} V_s \right) \cap V_{t^*} = V_t \cap V_{t^*}.$$

This proves that $x, y \in S(f)$, as desired. This concludes the proof. \blacksquare

Next, we define a family \mathcal{P}^* of k disjoint paths in H as follows. For every path $P \in \mathcal{P}$, let P^* denote the path in H that has vertex-set $V(P) \cap V_{t^*}$ and visits the vertices in $V(P) \cap V_{t^*}$ in the same order as P . This indeed forms a path in H , since for every subpath x, x_1, \dots, x_ℓ, y of P with $x, y \in V_{t^*}$ and $x_1, \dots, x_\ell \notin V_{t^*}$, we have $xy \in E(H)$ by definition.

For every endpoint v of a path $P \in \mathcal{P}$, let us denote by $v^* \in V_{t^*}$ the unique vertex in $V(P) \cap V_{t^*}$ that is closest to v along the path P . Necessarily, v^* is an endpoint of P^* . Using this notation, we now define two subsets $X^*, Y^* \subseteq V_{t^*}$ as

$$X^* := \{v^* : v \in X \text{ and } v \text{ is the endpoint of some path in } \mathcal{P}\},$$

$$Y^* := \{v^* : v \in Y \text{ and } v \text{ is the endpoint of some path in } \mathcal{P}\}.$$

In particular, \mathcal{P}^* is a collection of k pairwise disjoint X^* - Y^* -paths in H .

Claim 17. *There exists a family \mathcal{Q}^* consisting of pairwise non-adjacent X^* - Y^* -paths in H such that $|\mathcal{Q}^*| \geq \varepsilon k + \frac{1}{2}$.*

Proof of claim 17. By the properties of the tree-decomposition $(T, (V_t)_{t \in V(T)})$, we know that $\tau(t^*)$ either has at most r^2 vertices of degree at least $2r^4$, or is K_h -minor-free for $h = 2r^2$. In particular, the same is true for the subgraph H of $\tau(t^*)$.

Let us start with considering the first case. Let $Z \subseteq V_{t^*}$ denote the set of vertices of degree at least $2r^4$ in H . Then we know that $|Z| \leq r^2$ and $\Delta(H - Z) < 2r^4 := \Delta$. Let \mathcal{P}' be the set of paths of \mathcal{P}^* which do not intersect Z . Then, since the paths in \mathcal{P}^* are pairwise disjoint, we have that $|\mathcal{P}^*| \geq k - |Z| \geq k - 2r^4$. Hence, theorem 6 implies that there exists a collection \mathcal{Q}^* of pairwise non-adjacent $(X^* \setminus Z)$ - $(Y^* \setminus Z)$ -paths in $H - Z$ such that

$$|\mathcal{Q}^*| \geq \left\lfloor \frac{1}{2^{2\Delta^2}}(k - 2r^4) \right\rfloor.$$

Clearly, \mathcal{Q}^* is also a family of pairwise non-adjacent X^* - Y^* -paths in H . Using that $k > \frac{1}{2\varepsilon}$ and $\varepsilon = 2^{-(8t^8+3)}$, we can lower bound its size as follows.

$$\begin{aligned} |\mathcal{Q}^*| &\geq \frac{k}{2^{2\Delta^2}} - \frac{2r^4}{2^{2\Delta^2}} - 1 \\ &= \frac{k}{2^{8r^8}} - \frac{2r^4}{2^{8r^8}} - 1 \\ &= 8\varepsilon k - \frac{2r^4}{2^{8r^8}} - 1 \\ &= \varepsilon k + \frac{1}{2} + \left(7\varepsilon k - \frac{2r^4}{2^{8r^8}} - \frac{3}{2}\right) \\ &> \varepsilon k + \frac{1}{2}. \end{aligned}$$

This establishes the claim in the first case.

Next, consider the case that H is K_h -minor-free. Then, by Lemma 12 there exists a collection \mathcal{Q}^* of pairwise non-adjacent X^* - Y^* -paths in H of size at least $\frac{k}{2(h-1)} = \frac{k}{2(2r^2-1)} > 2^{-(8r^8+2)}k = 2\varepsilon k > \varepsilon k + \frac{1}{2}$, as desired. This concludes the proof of the claim also in the second possible case. \blacksquare

We now finish the proof of the theorem by using \mathcal{Q}^* , as given by claim 17, to construct a family \mathcal{Q} of pairwise non-adjacent X - Y -paths in G of size $|\mathcal{Q}| = |\mathcal{Q}^*| \geq \varepsilon k + \frac{1}{2}$.

For every edge $xy \in E(H) \setminus E(G[V_{t^*}])$, pick and fix a path P_{xy} in G that has endpoints x, y and no internal vertices in V_{t^*} (such a path always exists by definition of H). Furthermore, for every edge $xy \in E(G[V_{t^*}])$, we let P_{xy} denote the path consisting of the single edge xy in G .

Now consider any path $Q^* \in \mathcal{Q}^*$ and let x_1, x_2, \dots, x_q be its sequence of vertices, such that $x_1 \in X^*$ and $x_q \in Y^*$. Then, by definition, there exist $x \in X, y \in Y$ such that $x^* = x_1, y^* = x_q$. We now define $W(Q^*)$ as the walk⁵ in G that starts at x , follows the unique path in \mathcal{P} that x is an endpoint of, until it reaches $x^* = x_1$, then follows the

⁵Note that $W(Q^*)$ is not necessarily a path itself, given that the paths $P_{x_i x_{i+1}}$ might not be disjoint.

concatenation of the paths $P_{x_i x_{i+1}}$ for $1 \leq i < q$ and then follows the unique path in \mathcal{P} that y is an endpoint of, until it reaches y .

Claim 18. *If $Q_1^*, Q_2^* \in \mathcal{Q}^*$ are distinct, then $W(Q_1^*)$ and $W(Q_2^*)$ are non-adjacent in G .*

Proof of claim 18. Suppose towards a contradiction that there exist $a \in V(W(Q_1^*)), b \in V(W(Q_2^*))$ that are at distance at most 1 in G . Let $a' \in V(Q_1^*)$ be a vertex closest to a along $W(Q_1^*)$, and let $b' \in V(Q_2^*)$ be defined similarly for b . In particular, there exist paths R_1 and R_2 that form subwalks of $V(W(Q_1^*))$ and $V(W(Q_2^*))$, respectively, such that R_1 has endpoints a, a' and $V(R_1) \cap V_{t^*} = \{a'\}$, and analogously R_2 has endpoints b, b' and $V(R_2) \cap V_{t^*} = \{b'\}$. Now, as $a = b$ or $ab \in E(G)$, then the walk W in G that starts at a' , follows R_1 to a , moves to b , and follows R_2 until it reaches b' , satisfies $V(W) \cap V_{t^*} = \{a', b'\}$. This implies that there exists an a' - b' -path R in G with $V(R) \subseteq V(W)$, in particular we have $V(R) \cap V_{t^*} = \{a', b'\}$. By definition of H , this implies that $a' = b'$ or $a'b' \in E(H)$, in either case a contradiction, since Q_1^* and Q_2^* are non-adjacent in H . This concludes the proof of the claim. ■

We can now define \mathcal{Q} by, for every $Q^* \in \mathcal{Q}^*$, short-cutting the walk $W(Q^*)$ into a path in G that has the same endpoints and such that $V(Q) \subseteq V(W(Q^*))$. By claim 18, any two distinct paths in \mathcal{Q} are non-adjacent. Clearly, $|\mathcal{Q}| = |\mathcal{Q}^*|$ by definition, and claim 17 now implies that \mathcal{Q} consists of at least $\varepsilon k + \frac{1}{2}$ pairwise non-adjacent X - Y -paths in G . This yields the desired contradiction, completing the proof of the theorem. □

4 Subcubic graphs

In this section, we show our results on subcubic graphs. We begin by proving theorem 7, which we restate for convenience.

Theorem 7. *If G is a graph, $X, Y \subseteq V(G)$, and there exists a collection \mathcal{P} of at least $16k$ disjoint X - Y -paths such that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$, then there exist at least k pairwise non-adjacent X - Y -paths in G .*

Proof. Let $F := E(G[V(\mathcal{P})]) \setminus E(\mathcal{P})$. By theorem 10, it suffices to partition F into four induced matchings. By hypothesis, no two edges of F share an end vertex.

A standard tool for studying strong edge colouring is to find a proper vertex colouring of the square of the line graph; we slightly vary this argument given that we only want to partition the edges in F . We construct the auxiliary graph H with vertex set F and such that $e_1, e_2 \in F$ are adjacent if $\text{dist}_G(e_1, e_2) = 1$. In other words, e_1, e_2 are adjacent if and only if one of the end vertices of e_1 and one of the end vertices of e_2 are consecutive vertices on one of the paths in \mathcal{P} . This implies that $\Delta(H) \leq 4$, since the maximum degree in a path is two and since no two edges in F are incident.

By construction, a proper vertex-colouring of H with four colours yields the desired partition; each colour class is an induced matching. It suffices to show that H is 3-degenerate. Suppose, for a contradiction, that H has a 4-regular subgraph H' . Let $P \in \mathcal{P}$ such that some edge $e \in V(H')$ has an end $u \in V(P)$ and choose such e and u so




Figure 1: Example requiring four colours for any strong edge colouring of non-horizontal edges.

that the distance from u to an end of P along P is minimum. Then at most one neighbour of u in P is an end of an edge in $V(H')$. It follows that e has degree at most three in H' , obtaining the desired contradiction to the assumption that H' is 4-regular. \square

We note that the 4-colouring of the auxiliary graph H in the proof above is best possible. The configuration shown in fig. 1 (where the horizontal edges are part of the paths of \mathcal{P}), is an example in which we cannot partition the edges outside \mathcal{P} into three induced matchings. In particular, in this case the auxiliary graph is the Moser spindle [13], which is easily verified to not be 3-colourable.

We now prove theorem 8, which we restate for convenience.

Theorem 8. *If G is a graph, $X, Y \subseteq V(G)$, and there exists a collection \mathcal{P} of five disjoint X - Y -paths such that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$, then there exist two pairwise non-adjacent X - Y -paths in G .*

Furthermore, the statement does not necessarily hold if either

- (a) \mathcal{P} contains four paths instead of five, or
- (b) we replace the condition that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$ by the condition that the maximum degree of G is three.

Proof. We begin by proving the first part of the statement.

We define a *path system* as a quadruple $\mathcal{H} = (H, A, B, \mathcal{Q})$ where H is a graph, $A, B \subseteq V(H)$ and \mathcal{Q} is a 5-tuple of five disjoint A - B -paths such that

1. $V(H) = V(\mathcal{Q})$,
2. no vertex in A is incident in H to any edge in $E(H) \setminus E(\mathcal{Q})$,
3. every vertex of $V(H) \setminus A$ is incident in H to exactly one edge in $E(H) \setminus E(\mathcal{Q})$, and
4. there does not exist any collection \mathcal{Q}' of 5 pairwise disjoint A - B -paths such that $V(\mathcal{Q}') \subsetneq V(H)$.

Note that conditions (1) and (4) imply that the paths in \mathcal{Q} contain no vertices in $A \cup B$ other than their endpoints, and if $x \in A \cap B$ then one of the paths consists of exactly x . In particular $|A| = |B| = 5$.

We may suppose without loss of generality that $\mathcal{G} = (G, X, Y, \mathcal{P})$ is a path system. By restricting G , X and Y to vertices $V(\mathcal{P})$, we may suppose that (1) holds; any pair of

non-adjacent paths in an induced subgraph remains non-adjacent in the original graph. We may suppose (2) holds since if a vertex $x \in X$ is incident to some edge in $E(H) \setminus E(Q)$, we can add a new vertex x' to G as well as the edge $x'x$, replace X by $(X \setminus \{x\}) \cup \{x'\}$, and prepend $x'x$ to the path of \mathcal{P} with x as an endpoint. Any x' - Y -path in the new graph directly yields an x - Y -path with the same set of neighbours. We may suppose (3) holds for vertices not in Y given that finding a pair of non-adjacent paths in a graph directly yields such a pair in a subdivision of this graph. Furthermore, if some vertex $y \in Y \setminus X$ is not incident to some edge in $E(G) \setminus E(\mathcal{P})$, then y has a neighbour in the path of \mathcal{P} of which it is an endpoint, which exists by (1) (this path cannot be a singleton as $y \notin X$); say $y'y \in E(\mathcal{P})$. We may then remove y from G and from its path of \mathcal{P} and replace Y by $(Y \setminus \{y\}) \cup \{y'\}$; any X - y' -path directly extends to an X - y -path with the same neighbours. By repeating this argument, we may suppose that (3) holds for vertices in Y . Finally, we may suppose (4), as otherwise we could then replace \mathcal{P} with these paths. It is easily verified that none of these reductions are in conflict.

In the following, for $k \in \mathbb{N}$ we write $[k] = \{1, \dots, k\}$.

We now define an operation (which will have two variants) which allows us to easily construct and represent path systems. Let

$$\mathcal{H} = (H, A, B, \mathcal{Q} = (Q_1, Q_2, Q_3, Q_4, Q_5))$$

be a path system. For $i \in [5]$, write b_i for the vertex of B in Q_i .

Let $\mathcal{E} := \binom{[5]}{2}$ be the collection of (unordered) pairs of integers between 1 and 5. Let $\{i_1, i_2\} \in \mathcal{E}$. We define $\mathcal{H} \oplus \{i_1, i_2\}$ as the path system obtained by

- adding new vertices b'_{i_1}, b'_{i_2} and the edges $b_{i_1}b'_{i_1}, b_{i_2}b'_{i_2}$ and $b'_{i_1}b'_{i_2}$ to H ,
- appending the edges $b_{i_1}b'_{i_1}$ and $b_{i_2}b'_{i_2}$, respectively, to Q_{i_1} and Q_{i_2} , and
- replacing B with $(B \setminus \{b_{i_1}, b_{i_2}\}) \cup \{b'_{i_1}, b'_{i_2}\}$

An example of this operation is provided in fig. 2(b).

Let \mathcal{C} be the set of cyclic permutations of length at least two with values in $[5]$. We write such cycles as $(i_1 \dots i_k)$, for instance $(i_1 i_2 i_3) = (i_2 i_3 i_1) = (i_3 i_1 i_2)$. For $(i_1 \dots i_k) \in \mathcal{C}$, we say $\mathcal{H} \oplus (i_1 \dots i_k)$ is the path system obtained by

- adding new vertices c_{i_1}, \dots, c_{i_k} and $b'_{i_1}, \dots, b'_{i_k}$ and the edges $b_{i_1}c_{i_1}, \dots, b_{i_k}c_{i_k}, c_{i_1}b'_{i_1}, \dots, c_{i_k}b'_{i_k}$ and $c_{i_1}b'_{i_2}, \dots, c_{i_k}b'_{i_1}$ to H ,
- appending edges $b_{i_j}c_{i_j}$ and $c_{i_j}b'_{i_j}$ to Q_{i_j} for every $j \in [k]$, and
- replacing B with $(B \setminus \{b_{i_1}, \dots, b_{i_k}\}) \cup \{b'_{i_1}, \dots, b'_{i_k}\}$.

An example is provided in fig. 2(c).

Let $\mathcal{H}_0 = (H_0, A_0, B_0, \mathcal{Q}_0)$ be the path system consisting of the graph H_0 with singleton vertices $V(H_0) = A_0 = B_0 = \{v_1, v_2, v_3, v_4, v_5\}$ and \mathcal{P}_0 the 5 paths of length 0 in H_0 .




Figure 2: Example of a path system \mathcal{H} and two examples for the \oplus operation. The paths are labelled from 1 to 5 from top to bottom.

We now show that every path system can be obtained from \mathcal{H}_0 using the \oplus operation. We say two path systems $\mathcal{H}_1 = (H_1, A_1, B_1, \mathcal{Q}_1)$ and $\mathcal{H}_2 = (H_2, A_2, B_2, \mathcal{Q}_2)$ are *isomorphic*, which we denote by $\mathcal{H}_1 \simeq \mathcal{H}_2$, if there exists a graph isomorphism $h : V(H_1) \rightarrow V(H_2)$ which maps A_1 to A_2 , B_1 to B_2 and \mathcal{Q}_1 to \mathcal{Q}_2 (the ordering of the paths must be the same).

Claim 19. *For every path system \mathcal{H} , there exists some sequence $m_1, \dots, m_k \in \mathcal{E} \cup \mathcal{C}$ such that $\mathcal{H} \simeq \mathcal{H}_0 \oplus m_1 \oplus \dots \oplus m_k$.*

Proof of claim 19. Write $\mathcal{H} = (H, A, B, \mathcal{Q})$. We prove the statement by induction on $|E(H) \setminus E(\mathcal{Q})|$.

For the base case, if $|E(H) \setminus E(\mathcal{Q})| = 0$, then (3) implies that $A = V(\mathcal{Q})$, and so necessarily $\mathcal{H} \simeq \mathcal{H}_0$.

We now show the inductive step. Let F be the set of edges in $E(H) \setminus E(\mathcal{Q})$ incident to B . As $|E(H) \setminus E(\mathcal{Q})| > 0$, condition (2) implies that $A \neq B$, and so by (3) we have that $F \neq \emptyset$.

First suppose there exists an edge of F with both ends in B , say $b'_{i_1}b'_{i_2}$, where $b'_{i_1} \in V(Q_{i_1})$ and $b'_{i_2} \in V(Q_{i_2})$. Let \mathcal{H}' be the path system resulting from removing b'_{i_1}, b'_{i_2} . More precisely, if b_{i_1} and b_{i_2} are, respectively, the neighbours of b'_{i_1} in Q_{i_1} and of b'_{i_2} in Q_{i_2} (these exist since $b'_{i_1}, b'_{i_2} \notin A$), then $\mathcal{H}' = (H', A, B', \mathcal{Q}')$ where $H' = H - \{b'_{i_1}, b'_{i_2}\}$, $B' = (B \setminus \{b'_{i_1}, b'_{i_2}\}) \cup \{b_{i_1}, b_{i_2}\}$ and \mathcal{Q}' is identical to \mathcal{Q} except that the edges $b_{i_1}b'_{i_1}$ and $b_{i_2}b'_{i_2}$ are removed from, respectively, Q_{i_1} and Q_{i_2} . It is direct from the definitions that $\mathcal{H} \simeq \mathcal{H}' \oplus \{i_1, i_2\}$. In particular, condition (3) implies that $b'_{i_1}b'_{i_2}$ were not incident to any edge other than $b_{i_1}b'_{i_1}$, $b_{i_2}b'_{i_2}$ and $b'_{i_1}b'_{i_2}$. Furthermore, $E(H') \setminus E(\mathcal{Q}') = (E(H) \setminus E(\mathcal{Q})) \setminus \{b'_{i_1}b'_{i_2}\}$. By induction, there exists a sequence $m_1, \dots, m_{k-1} \in \mathcal{E} \cup \mathcal{C}$ such that $\mathcal{H}' \simeq \mathcal{H}_0 \oplus m_1 \oplus \dots \oplus m_{k-1}$. Hence, we obtain that $\mathcal{H} \simeq \mathcal{H}_0 \oplus m_1 \oplus \dots \oplus m_{k-1} \oplus \{i_1, i_2\}$, as desired.

Otherwise, we construct an auxiliary digraph J with vertex set F as follows. Let $c_{i_1}b'_{i_2}, c_{i_3}b'_{i_4} \in F$, where $b'_{i_2} \in V(Q_{i_2}) \cap B$ and $b'_{i_4} \in V(Q_{i_4}) \cap B$, and $c_{i_1} \in V(Q_{i_1}) \setminus B$ and $c_{i_3} \in V(Q_{i_3}) \setminus B$. Further note that $i_1 \neq i_2$ and $i_3 \neq i_4$: these edges are not in \mathcal{Q} and by (4) the paths of \mathcal{Q} are necessarily induced. In our auxiliary digraph J , we put a directed edge from $c_{i_1}b_{i_2}$ to $c_{i_3}b_{i_4}$ in J if and only if $i_2 = i_3$.

Every edge in F has in-degree at least one in J (in fact, it is necessarily exactly one). Indeed, let $c_{i_1}b'_{i_2} \in F$. Let b'_{i_1} be the unique vertex of $V(Q_{i_1}) \cap B$. As $c_{i_1} \notin B$, necessarily $b'_{i_1} \notin A$. Hence, by (3), there exists there is some edge $e \in F$ incident to b'_{i_1} . In particular, there is a directed edge in J from e to $c_{i_1}b'_{i_2}$.

Hence, there exists in J a directed cycle. By definition, the sequence of vertices in this directed cycle is of the form $c_{i_1}b'_{i_2}, \dots, c_{i_k}b'_{i_1}$, where $c_{i_j} \in V(Q_{i_j}) \setminus B$ and $b'_{i_j} \in V(Q_{i_j}) \cap B$ for every $j \in [k]$. We claim $c_{i_j}b'_{i_j} \in E(Q_{i_j})$ for every $j \in [k]$. Suppose otherwise that for some $j \in [k]$ there exists at least one vertex x between c_{i_j} and b_{i_j} on Q_{i_j} . Then, the following collection of paths would contradict (4): in \mathcal{Q} , replace the paths Q_{i_1}, \dots, Q_{i_k} with the paths formed by following P_{i_j} until c_{i_j} and then following the edge $c_{i_j}b'_{i_{j+1}}$ (with addition modulo k). In particular, this new set of paths does not contain x .

Let \mathcal{H}' be the path system resulting from removing c_{i_j} and b'_{i_j} for every $j \in [k]$. More precisely, if we write b_{i_j} for the neighbour of c_{i_j} in $V(Q_{i_j})$ which is not b'_{i_j} (this vertex exists since $c_{i_j} \notin A$ by (2)), for every $j \in [k]$, then $\mathcal{H}' = (H', A, B', \mathcal{Q}')$ where $H' = H - \{c_{i_1}, \dots, c_{i_k}, b'_{i_1}, \dots, b'_{i_k}\}$, $B' = (B \setminus \{b'_{i_1}, \dots, b'_{i_k}\}) \cup \{b_{i_1}, \dots, b_{i_k}\}$ and \mathcal{Q}' is identical to \mathcal{Q} except that the edges $b_{i_j}c_{i_j}$ and $c_{i_j}b'_{i_j}$ are removed from Q_{i_j} , for every $j \in [k]$. Similarly to above, it is then direct from the definitions that $\mathcal{H} \simeq \mathcal{H}' \oplus (i_1 \dots i_k)$, and by induction, there exists a sequence $m_1, \dots, m_{k-1} \in \mathcal{E} \cup \mathcal{C}$ such that $\mathcal{H}' \simeq \mathcal{H}_0 \oplus m_1 \oplus \dots \oplus m_{k-1}$. Hence, we obtain that $\mathcal{H} \simeq \mathcal{H}_0 \oplus m_1 \oplus \dots \oplus m_{k-1} \oplus (i_1 \dots i_k)$, as desired. \blacksquare

We define a *state* as an unordered pair $\{S_1, S_2\}$ of non-empty disjoint subsets of $[5]$. Given a path system $\mathcal{H} = (H, A, B, \mathcal{Q} = (Q_1, Q_2, Q_3, Q_4, Q_5))$, we say a state $S = \{S_1, S_2\}$ is \mathcal{H} -*reachable* if there exist sets $C_1, C_2 \subseteq V(H)$ such that

- C_1, C_2 are disjoint and non-adjacent in H ,
- $C_1 \cap A \neq \emptyset$ and $C_2 \cap A \neq \emptyset$,
- $C_1 \cap V(Q_i) \cap B \neq \emptyset$ if and only if $i \in S_1$ and $C_2 \cap V(Q_i) \cap B \neq \emptyset$ if and only if $i \in S_2$, and
- $H[C_1], H[C_2]$ are connected.

Claim 20. *For every path system \mathcal{H} , there exists an \mathcal{H} -reachable state.*

Proof of claim 20. Let $S = \{S_1, S_2\}$ be a state and $m \in \mathcal{E} \cup \mathcal{C}$. We construct a collection of states $f(S, m)$ by saying $S' \in f(S, m)$ if and only if S' can be written as $S' = \{S_1, S_2\}$ such that the following holds: writing, $\mathcal{H}_0 \oplus m = (H_m, A_m, B_m, \mathcal{Q}_m)$, there exist $C_1, C_2 \subseteq V(H_m)$ such that

- C_1, C_2 are disjoint and non-adjacent in H_m ,
- $C_1 \cap V(Q_i) \cap A_m$ if and only if $i \in S_1$ and $C_2 \cap V(Q_i) \cap A_m$ if and only if $i \in S_2$,
- $C_1 \cap V(Q_i) \cap B_m \neq \emptyset$ if and only if $i \in S'_1$ and $C_2 \cap V(Q_i) \cap B_m \neq \emptyset$ if and only if $i \in S'_2$,

- for every vertex of $C_1 \cap B_m$, its connected component in $H[C_1]$ contains a vertex of $C_1 \cap A_m$, and for every vertex of $C_2 \cap B_m$, its connected component in $H[C_2]$ contains a vertex of $C_2 \cap A_m$.

It is easily verified that for any path system \mathcal{H} , if S is \mathcal{H} -reachable, then every state in $f(S, m)$ is $\mathcal{H} \oplus m$ -reachable. The crucial observation is that $\mathcal{H} \oplus m$ can be obtained from \mathcal{H} and $\mathcal{H}_0 \oplus m$ by identifying the vertices in the B set from the former and the vertices from the A set in the latter. From this, these conditions are exactly those which allow us to extend the non-adjacent sets C_1, C_2 corresponding to states in \mathcal{H} to sets in $\mathcal{H} \oplus m$.

If \mathcal{S} is a collection of states, let $g(\mathcal{S}, m) = \bigcup_{S \in \mathcal{S}} f(S, m)$. Hence, for any \mathcal{H} , if \mathcal{S} is a collection of \mathcal{H} -reachable states, then $g(\mathcal{S}, m)$ is a collection of $\mathcal{H} \oplus m$ -reachable states. We say a collection of states \mathcal{S}' is a *descendant* of \mathcal{S} if there exists some sequence $m_1, \dots, m_k \in \mathcal{E} \cup \mathcal{C}$ such that $\mathcal{S}' = g(\dots g(g(\mathcal{S}, m_1), m_2) \dots, m_k)$.

Let $\mathcal{S}_0 := \{\{\{x\}, \{y\}\} : x, y \in [5], x \neq y\}$. It is direct from the definition that every state in \mathcal{S}_0 is \mathcal{H}_0 -reachable. By claim 19, every \mathcal{H} can be written as $\mathcal{H} = \mathcal{G}_0 \oplus m_1 \oplus \dots \oplus m_k$ for some sequence of $m_1, \dots, m_k \in \mathcal{E} \cup \mathcal{C}$. In particular, $g(\dots g(g(\mathcal{S}_0, m_1), m_2) \dots, m_k)$ is a collection of \mathcal{H} -reachable states. We want to show that this collection is non-empty. In order to prove the claim, it thus suffices to show that \emptyset is not a descendant of \mathcal{S}_0 . Our strategy is thus as follows : start with the collection \mathcal{S}_0 , and then repeatedly apply $g(\cdot, m)$ for every $m \in \mathcal{E} \cup \mathcal{C}$ until no new collection of states are found. If the \emptyset is never encountered, we are done.

Along the way, we may in fact trim some branches off of this process, in the three following ways:

- (i) If \mathcal{S} is a collection of \mathcal{H} -reachable states and σ is a permutation of $[5]$, then $\sigma(\mathcal{S})$, which is obtained by applying σ to the elements in \mathcal{S} , is \mathcal{H}' -reachable, where \mathcal{H}' is identical to \mathcal{H} except that we have permuted the order of the paths of \mathcal{Q} according to σ .

If $g(\dots g(g(\sigma(\mathcal{S}), m_1), m_2) \dots, m_k) = \emptyset$, then

$$g(\dots g(g(\mathcal{S}', \sigma^{-1}(m_1)), \sigma^{-1}(m_2)) \dots, \sigma^{-1}(m_k))$$

is empty, and so if \emptyset is not a descendant of \mathcal{S} , it is not a descendant of $\sigma(\mathcal{S})$ either. In particular, we need only keep one collection of states from each equivalence class under permutation.

- (ii) If $\mathcal{S} \subseteq \mathcal{S}'$ are collections of states, then $g(\mathcal{S}, m) \subseteq g(\mathcal{S}', m)$. In particular, if \emptyset is not a descendant of \mathcal{S} , it is not a descendant of \mathcal{S}' either. Hence, during our searching process, we may throw out any collection of states which is not minimal.
- (iii) We do not need to use the definition of f to compute every instance of $f(S, m)$. Indeed, similarly to (i), it is easily verified from the definitions that if σ is a permutation of $[5]$, then $f(S, m) = \sigma^{-1}(f(\sigma(S), \sigma(m)))$, so we only need to compute f for one pair (S, m) in every equivalence class.

We have implemented this approach in Mathematica [15], the code is provided in [8]. This code is fully commented, consult these for further implementation details. As a benchmark, this script can run in under 14 minutes on a 2020 MacBook Air with M1 chip and 16 GB ram running Mathematica 13.0.0.0. ■

In particular, there exists some \mathcal{G} -reachable state. This state is a certificate of the existence of two non-adjacent sets C_1, C_2 which induce connected graphs and both intersect X and Y . From these, we may extract two non-adjacent X - Y -paths, as desired. This concludes the proof of the first part of the statement.

We now show that this result is best possible, in two ways.

First, consider the statement when we are given four paths instead of five. By modifying the approach used above to also take into account paths which “backtrack”, we were able to find a counter-example to this modified statement; it is shown in fig. 3. A short Mathematica script which verifies that no pair of non-adjacent X - Y -paths exists in this graph is provided in [8].




Figure 3: Counter-example (a) in theorem 8.




Figure 4: Counter-example (b) in theorem 8.

Secondly, consider the statement if we replace the condition that every vertex in $V(\mathcal{P})$ is incident to at most one edge in $E(G) \setminus E(\mathcal{P})$ by the condition that the maximum degree is three. Of course, the difference only concerns vertices in X and in Y , given that these vertices are the only ones with fewer than two neighbours in their path. The graph formed by adding a matching between a five-cycle (with vertices X) and the complement of a copy of this five-cycle (with vertices Y), as shown in fig. 4, is easily verified to not contain any pair of non-adjacent edges between X and Y . □

Acknowledgements

The first author is supported by the Institute for Basic Science (IBS-R029-C1). The second and fourth authors are supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). Les deuxième et quatrième auteurs sont supportés par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG). The third author is funded by an ETH Zürich Postdoctoral Fellowship.

References

- [1] S. Albrechtsen, T. Huynh, R. W. Jacobs, P. Knappe, and P. Wollan. The induced two paths problem, June 2023. [arXiv:2305.04721](https://arxiv.org/abs/2305.04721).
- [2] L. D. Andersen. The strong chromatic index of a cubic graph is at most 10. *Discrete Mathematics*, 108(1-3):231–252, Oct. 1992.
- [3] P. Duchet and H. Meyniel. On Hadwiger’s Number and the Stability Number. *North-Holland Mathematics Studies*, 62:71–73, 1982.
- [4] J. Erde and D. Weißauer. A Short Derivation of the Structure Theorem for Graphs with Excluded Topological Minors. *SIAM Journal on Discrete Mathematics*, 33(3):1654–1661, Jan. 2019.
- [5] P. Gartland, T. Korhonen, and D. Lokshitanov. On Induced Versions of Menger’s Theorem on Sparse Graphs, Sept. 2023. [arXiv:2309.08169](https://arxiv.org/abs/2309.08169).
- [6] A. Georgakopoulos and P. Papasoglu. Graph minors and metric spaces, Oct. 2023. [arXiv:2305.07456v2](https://arxiv.org/abs/2305.07456v2).
- [7] M. Grohe and D. Marx. Structure Theorem and Isomorphism Test for Graphs with Excluded Topological Subgraphs. *SIAM Journal on Computing*, 44(1):114–159, Jan. 2015.
- [8] K. Hendrey, S. Norin, R. Steiner, and J. Turcotte. Code used in the paper “An induced version of Menger’s theorem”, 2023. <https://github.com/tjeremie/Induced-Menger>.
- [9] P. Horák, H. Qing, and W. T. Trotter. Induced matchings in cubic graphs. *Journal of Graph Theory*, 17(2):151–160, June 1993.
- [10] E. Hurley, R. d. J. de Verclos, and R. J. Kang. An improved procedure for colouring graphs of bounded local density. *Advances in Combinatorics*, Sept. 2022:7.
- [11] T. Korhonen. Grid induced minor theorem for graphs of small degree. *Journal of Combinatorial Theory, Series B*, 160:206–214, May 2023.
- [12] K. Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927.
- [13] L. Moser and W. Moser. Problems for Solution. *Canadian Mathematical Bulletin*, 4(2):187–187, May 1961.

- [14] T. Nguyen, A. Scott, and P. Seymour. A counterexample to the coarse Menger conjecture, Jan. 2024. [arXiv:2401.06685](https://arxiv.org/abs/2401.06685).
- [15] Wolfram Research Inc. Mathematica. <https://www.wolfram.com/mathematica/>.

A Reduction of conjecture 4 to distance 3

In this appendix, we show that reduction for $d \geq 3$ to $d = 3$ by McCarty and Seymour also holds for conjecture 4. For this paper to be self-contained, we reproduce this proof here with only slight modifications for the change of setting.

Theorem 21. *If conjecture 4 holds for $d = 3$, it also holds for $d \geq 3$ with constant $C(d, \Delta) = C(3, \Delta^{d+1})k$.*

Proof. Suppose G be a graph with $\Delta(G) \leq \Delta$ and $X, Y \subseteq V(G)$, and $k \in \mathbb{N}$. Let H be the d -th power of G , i.e. the graph obtained from G by adding edges to any pair of vertices at distance at most d in G . Note that the maximum degree of H is at most Δ^{d+1} . Given that conjecture 4 holds for distance 3, there exist either

- (1) k X - Y -paths pairwise at distance at least 3 in H , or
- (2) a set of fewer than $C(3, \Delta^{d+1})k$ vertices in H which separates X and Y .

First suppose we are in case (1). Let P'_1, \dots, P'_k be X - Y paths in H , pairwise at distance at least 3. We wish to construct P_1, \dots, P_k , X - Y paths in H , pairwise at distance at least d . For each $i \in [k]$, P'_i can be converted into a path in G with the same endpoints by following paths of length at most d in G between every pair of consecutive vertices in P'_i (these exist by definition of H), and then converting the resulting walk to a path by removing cycles. In particular, any vertex in P_i is at distance at most $\frac{d}{2}$ from P'_i . If $d_G(P_i, P_j) \geq d$ for every distinct $i, j \in [k]$, then we are done. Otherwise, suppose there exist distinct $i, j \in [k]$ such that $\text{dist}_G(P_i, P_j) \leq d$. In particular, the distance between P'_i and P'_j is at most $\frac{d}{2} + d + \frac{d}{2} = 2d$. By definition of H , this implies that $\text{dist}_H(P'_i, P'_j) \leq 2$, which is a contradiction.

Otherwise, we are in case (2). As G is a subgraph of H , the set of size less than $C(3, \Delta^{d+1})k$ which separates X and Y in H also separates X and Y in G . \square