Long cycles have the edge-Erdős-Pósa property

Henning Bruhn, Matthias Heinlein and Felix Joos*

Abstract

We prove that the set of long cycles has the edge-Erdős-Pósa property: for every fixed integer $\ell \geq 3$ and every $k \in \mathbb{N}$, every graph G either contains k edge-disjoint cycles of length at least ℓ (long cycles) or an edge set X of size $O(k^2 \log k + \ell k)$ such that $G - X$ does not contain any long cycle. This answers a question of Birmelé, Bondy, and Reed (Combinatorica 27 (2007), 135–145).

1 Introduction

Many theorems have a vertex version and an edge version. There is a Menger theorem about (vertex-)disjoint paths and a variant about edge-disjoint paths. We prove here the edge analogue of an Erdős-Pósa-type theorem.

Erdős and Pósa [6] proved in 1962 that every graph either contains k disjoint cycles or a set of $O(k \log k)$ vertices that meets every cycle. Since then many Erdős-Pósa-type theorems have been discovered, among them one about *long* cycles. These are cycles of a length that is at least some fixed integer ℓ.

Indeed, every graph either contains k disjoint long cycles or a set of $O(k\ell + k \log k)$ vertices that meets every cycle. With a worse bound this follows from a theorem of Robertson and Seymour [18], while the stated bound is due to Mousset, Noever, Škorić, and Weissenberger [14]. We prove an edge-disjoint analogue:

Theorem 1. *Let ℓ be a positive integer. Then every graph G either contains k edge-disjoint long cycles or a set $X \subseteq E(G)$ of size $O(k\ell + k^2 \log k)$ such that $G - X$ contains no long cycle.*

This answers a question of Birmelé, Bondy, and Reed [2].

For vertex-disjoint long cycles, the bound of $O(k\ell + k \log k)$ proved by Mousset et al. [14] is optimal as it matches a lower bound found by Fiorini and Herinckx [7]. We show below that the set X in Theorem 1 also needs to have size at least $\Omega(k\ell + k \log k)$. We believe that, as in the vertex version, this is the right order of magnitude.

A family $\mathcal{F}$ of graphs has the *Erdős-Pósa property* if there is a function $f_{\mathcal{F}} : \mathbb{N} \to \mathbb{R}$ such that for every integer k every graph G either contains k disjoint copies of graphs in $\mathcal{F}$ or a *hitting set* $X \subseteq V(G)$ of size at most $f_{\mathcal{F}}(k)$ that meets every $\mathcal{F}$-copy in G. Thus cycles have the Erdős-Pósa property, but also, for instance, even cycles [20] and many other graph classes.

*The research was also supported by the EPSRC, grant no. EP/M009408/1.

Many such results are the consequence of a far-reaching theorem of Robertson and Seymour [18]: for a fixed graph H the class of graphs that have H as a minor has the Erdős-Pósa property if and only if H is planar. For example, the theorem implies that long cycles have the Erdős-Pósa property.

Less is known about the edge analogue of the Erdős-Pósa property. There, the objective is to find edge-disjoint copies of graphs in $\mathcal{F}$ or a bounded hitting set of edges. While cycles have the *edge-Erdős-Pósa property* [5, Exercise 9.5], an edge version of Robertson and Seymour's theorem, for example, is still wide open. By our result, long cycles have the edge-Erdős-Pósa property.

We know of only two other graph classes that have the edge-Erdős-Pósa property: S-cycles, cycles that each contain a vertex from a fixed set S, and the graphs that contain a θ_r-minor, where θ_r is the multigraph consisting of two vertices linked by r parallel edges. The first result is due to Pontecorvi and Wollan [15], the second due to Raymond, Sau and Thilikos [16]. Strikingly, both results are obtained via a reduction to their respective vertex versions. For long cycles this does not seem to be possible (at least not that easily), and consequently, our proof is direct.

Within restricted ambient graphs, two more graph classes are known to have the edge-Erdős-Pósa property. Odd cycles do not have the Erdős-Pósa property, and they do not have the edge version either [4]. The same is true for the class of graphs that contain an immersion[1] of H for certain graphs H. If, however, the ambient graphs G are required to be 4-edge-connected, then odd cycles as well as graphs with an H-immersion gain the edge-Erdős-Pósa property [10, 11].

There are many more results about the ordinary Erdős-Pósa property, most of which are listed in the survey of Raymond and Thilikos [17]. A direction we find interesting concerns rooted graphs. In this setting, a set S (or two or more such sets) is fixed in the ambient graph G. The target objects are required to meet the set S in some specified way. For instance, S-cycles, cycles that each intersect S, have the Erdős-Pósa property [9, 15], and this is still true for long S-cycles [3]. Huynh, Joos, and Wollan [8] verify the Erdős-Pósa property for cycles satisfying more general restrictions that include for example S_1-S_2-cycles (defined in the obvious way). Note that S_1-S_2-S_3-cycles do not have the Erdős-Pósa property. We do not know whether the Erdős-Pósa property extends to edge-disjoint S_1-S_2-cycles.

In Section 2, we discuss the size of the hitting set and how the Erdős-Pósa property and its edge analogue differ. In Section 3, we introduce tools needed in the proof of Theorem 1. After a brief overview we prove Theorem 1 in Section 4.

2 Discussion

2.1 The size of the hitting set

Fiorini and Herinckx [7] observed that the hitting set for long cycles in the ordinary Erdős-Pósa property needs to have size at least $\Omega(k\ell + k \log k)$. That there is a hitting set of size $O(k\ell + k \log k)$, the optimal size, is due to Mousset

[1] A graph G contains an immersion of H if there is an injective function $\tau : V(H) \to V(G)$ and edge-disjoint $\tau(u)$–$\tau(v)$-paths for every $uv \in E(H)$ in G.

et al. [14] who built on earlier work of Robertson and Seymour [18], Birmelé et al. [2], and Fiorini and Herinckx [7].

What is the optimal size of the hitting set in the edge-disjoint version? As for vertex-disjoint long cycles, the construction of Simonovits [19], originally intended for the classic Erdős-Pósa theorem, gives a lower bound of $\Omega(k \log k)$. Indeed, the graphs in the construction are cubic, which means that cycles are disjoint if and only if they are edge-disjoint.

That the size of the hitting set needs to depend on ℓ at all is not immediately obvious. But it does, and indeed, the dependence is linear. To prove this we construct graphs S_ℓ that do not contain two edge-disjoint long cycles and that do not admit a hitting set of less than $\frac{\ell}{30}$ edges. Taking $k-1$ disjoint copies of S_ℓ then yields a graph without k edge-disjoint cycles and no hitting set of size smaller than $\frac{1}{30}(k-1)\ell = \Omega(k\ell)$. Therefore, the size of hitting sets for edge-disjoint long cycles needs to be at least $\Omega(k\ell + k \log k)$.

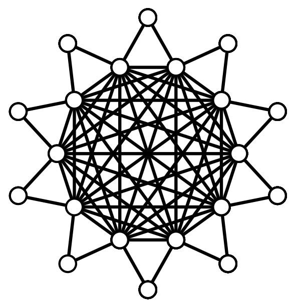

Figure 1: The graph S_{17} contains no two edge-disjoint cycles of length at least 17.

The graphs S_ℓ are constructed as follows. Let $p = \lfloor \frac{2}{3}(\ell - 1) \rfloor$, and let S_ℓ be the graph obtained from a clique on p vertices $v_0, \ldots, v_{p-1}$ by adding vertices $w_0, \ldots, w_{p-1}$ such that each w_i is adjacent to v_{i-1} and v_i (where we take indices mod p). The graphs S_ℓ are sometimes called *suns* [1]. As the clique contains only $p < \frac{2}{3}\ell$ vertices, every long cycle in S_ℓ passes through at least $\frac{1}{3}\ell + 1 \geq \frac{p}{2} + 1$ vertices of $\{w_0, \ldots, w_{p-1}\}$. As these have degree 2, there cannot be two edge-disjoint long cycles in S_ℓ.

Let $\ell \geq 30$, and consider any set X of at most $\frac{\ell}{30}$ edges. We show that X is not a hitting set. For every edge $uv \in X$ delete its endvertices u and v in G, and if we delete a vertex v_i of the clique, also delete the adjacent vertices w_i and w_{i+1}. All in all, we delete a set U of at most $6 \cdot \frac{\ell}{30} \leq \frac{\ell}{5}$ vertices in G. For the cycle $C = v_0 \ldots v_{p-1} v_0$, let $C_1, \ldots, C_r$ be the components of $C - U$. Let v_{s_i} and v_{t_i} be the two endpoints of the paths C_i. None of the vertices $w_{s_i}, \ldots, w_{t_i - 1}$ is deleted, and thus $P_i = v_{s_i} w_{s_i} v_{s_i+1} \ldots w_{t_i-1} v_{t_i}$ is a path in $G - U$.

Concatenating the paths P_i by adding the edges $v_{t_i} v_{s_{i+1}}$, we obtain a Hamilton cycle D of $G - U$. Noting that $p \geq \frac{2}{3}(\ell - 3)$, we calculate that the length of D is

$$|V(S_\ell)| - |U| = 2p - \frac{1}{5}\ell \geq \frac{4}{3}(\ell - 3) - \frac{1}{5}\ell = \ell + \frac{2\ell - 60}{15} \geq \ell$$

as $\ell \geq 30$. Since $G - X \supseteq G - U$ still contains a long cycle, we deduce that no edge set of size at most $\frac{\ell}{30}$ is a hitting set.

Comparing the lower bound of $\Omega(k\ell + k\log k)$ with Theorem 1, we see that there is a gap in the second term by a factor k. We believe that the optimal size of the hitting set coincides with the lower bound. In one argument our proof seems to be wasteful by an additional factor of k. Unfortunately, we have been unable to do the step in a more economical way.

2.2 Vertex versus edge version

Why is the edge-Erdős-Pósa property hard at all, especially when the corresponding vertex version is known? Cannot a reduction be employed or the proof be adapted? Pontecorvi and Wollan [15] obtain the edge version for S-cycles from the vertex version by a simple gadget construction. Essentially, they apply the vertex version to a modified line graph (a similar approach is also used by Kawarabayashi and Kobayashi [10]). Why is that not possible for long cycles?

Cycles do not have a unique image in the line graph. The line graph of a cycle is a cycle but not every cycle in the line graph corresponds to a cycle in the root graph. The preimage of an S-cycle in the (slightly modified) line graph still contains an S-cycle—this is what allows Pontecorvi and Wollan to reduce to the vertex version. For long cycles this will not work because every cycle contained in the preimage of a long cycle might be short.

So how about adapting the proof of the vertex version in some more or less obvious way? While the existing proof might, and does in our case, give some clues, an easy adaption seems hopeless. We believe this is because edge-disjoint long cycles actually require a mix of the two disjointness concepts.

Why is this? For simplicity, consider the case $k = 2$. We could construct two long cycles in a graph G as follows. Choose 2ℓ vertices $v_1, \dots, v_\ell$ and $w_1, \dots, w_\ell$. For the vertex version, suppose that all these vertices are distinct. What we now need to do is to find internally vertex-disjoint paths $P_1, \dots, P_\ell$ and $Q_1, \dots, Q_\ell$ such that P_i is a v_i–v_{i+1}-path and Q_i a w_i–w_{i+1}-path for every $i = 1, \dots, \ell$ (where we set $v_{\ell+1} = v_1$ and $w_{\ell+1} = w_1$). In the edge version, we only need to suppose that $v_i \neq v_j$ and $w_i \neq w_j$ for distinct i, j. Again, we seek for paths connecting these vertices in cyclic order. But, and that is the crucial point, P_i and P_j as well as Q_i and Q_j need to be internally vertex-disjoint for distinct i, j, while P_i and Q_j only need to be edge-disjoint. That is, we deal with two different types of disjointness.

If instead we only require that all these paths are edge-disjoint, then we obtain immersions of long cycles. Strikingly, for immersions the adaption of vertex version arguments appears to work very well. Indeed, to prove his strong result about edge-disjoint immersions, Liu [11] translates a part of the graph minor theory to line graphs. (The translation, however, is not at all trivial.)

3 Preliminaries

In this section we introduce some notation. In particular, we define extensions of paths and frames. We devote to each of these concepts a subsection where we collect a few properties about these. In another subsection, we prove several results about edge-connected multigraphs.

All logarithms $\log n$ will be to base 2.

3.1 Paths and cycles

We follow the notation used in the textbook of Diestel [5]. In particular, we write $P = u \dots v$ for a path P with endvertices u, v and say that P is a *u–v-path*. For two vertices $x, y \in V(P)$, we denote by xPy the subpath of P with endvertices x, y. We also write xCy for an oriented cycle C and $x, y \in V(C)$ to denote the x–y-subpath of C. For paths $x_1P_1y_1$ and $x_2P_2y_2$ such that $x_2 = y_1$ and otherwise P_1 and P_2 are disjoint, we write $x_1P_1x_2P_2y_2$ for the concatenation of P_1 and P_2. For two vertex sets A, B, we define an *A–B-path* as a path P such that one endpoint of P lies in A and one in B and P is internally disjoint from $A \cup B$. For a subgraph H of G (or a vertex set which we treat as a subgraph without edges), we define an *H-path* is a path with endvertices in H that is internally disjoint from H. Note that the path may have length 1.

For a cycle C and a path P, we denote by $\ell(C)$ and $\ell(P)$ the number of edges of C and P, respectively, and refer to $\ell(C)$ and $\ell(P)$ as the *length* of C and P, respectively.

Throughout the article, we fix a positive integer ℓ and call P and C *short* if $\ell(P) < \ell$ and $\ell(C) < \ell$, respectively. A cycle is called *long* if its length is at least ℓ.

3.2 Extensions of paths

The key trick in our proofs is to exclude cycles of intermediate length, that is, cycles that are long but not too long. In this subsection we treat a tool, path extensions, that allows us to construct such intermediate cycles. Since these are excluded we will then obtain the desired contradiction.

Consider a path P with endvertices u, v. We write $\leq_P$ for the total order of the vertices $V(P)$ induced by the distance from u on P. Let $Q_1, \dots, Q_r$ be P-paths, and for $i = 1, \dots, r$ let u_i and v_i be the endvertices of Q_i such that $u_i <_P v_i$. The tuple $(Q_1, \dots, Q_r)$ is an *extension* of P if

(E1) the paths $Q_1, \dots, Q_r$ are pairwise internally disjoint;

(E2) the cycle $u_iPv_i \cup Q_i$ is short for $i = 1, \dots, r$;

(E3) $u_1 = u$ and $v_r = v$;

(E4) $u_i <_P u_{i+1} <_P v_i <_P v_{i+1}$ for $i = 1, \dots, r-1$; and

(E5) $v_i \leq_P u_{i+2}$ for $i = 1, \dots, r-2$.

See Figure 2 for an illustration.

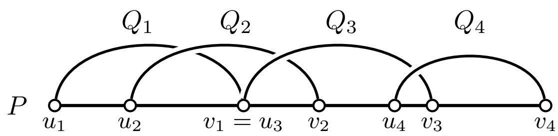

Figure 2: A P-extension.

Lemma 2. *Let P be a path, and let $(Q_1, \dots, Q_r)$ be an extension of P. For any i, j with $1 \leq i \leq j \leq r$, there is exactly one cycle C in $P \cup \bigcup_{s=i}^{j} Q_s$ that contains u_i, v_j. The edge set of the cycle is*

$$E(C) = E\big(P \cup \bigcup_{s=i}^{j} Q_s\big) \setminus \bigcup_{t=i+1}^{j} E(u_t P v_{t-1}) \tag{1}$$

Proof. The graph $H = P \cup \bigcup_{s=i}^{j} Q_s$ is 2-connected as it is the union of cycles $u_s P v_s \cup Q_s$, such that consecutive cycles overlap in an edge. Thus the graph H contains a cycle C through u_i and v_j.

Note that C has to contain each of $Q_i, \dots, Q_{j-1}$: if $Q_t \nsubseteq C$ for a $t \in \{i, \dots, j-1\}$ then, by (E5), u_{t+1} separates u_i and v_j in C , which is impossible. We also have $Q_j \subseteq C$ as otherwise v_j would have degree 1 in C as $v_j \notin Q_{j-1}$ by (E4).

Now, for $t = i+1, \dots, j$ the vertex v_{t-1} has degree 2 in C. Therefore, either $u_t P v_{t-1} \subseteq C$ or $v_{t-1} P u_{t+1} \subseteq C$ (where we temporarily interpret u_{j+1} as v_j). However, $\{u_t, v_{t-1}\}$ separates u_i from v_j in H, which means that C has to pass through $\{u_t, v_{t-1}\}$ twice. Thus $v_{t-1} P u_{t+1} \subseteq C$ and $u_t P v_{t-1} \nsubseteq C$ (since already $Q_{t-1} \subseteq C$). It is easy to check that this fixes C to be as in (1). □

Lemma 3. *Let P be a path, and let $(Q_1, \dots, Q_r)$ be an extension of P. Assume that every long cycle in $H = P \cup \bigcup_{s=1}^{r} Q_s$ has length at least 2ℓ. Then every cycle in H is short.*

Proof. Suppose that H contains a long cycle C. Clearly, its intersection with P is nonempty. Let i be the smallest index such that u_i lies in C, and let j be the largest index with $v_j \in V(C)$. Note that $i < j$ by the definition of extensions. We, furthermore, assume C to be chosen such that $j - i$ is minimal. Thus $C \subseteq u_i P v_j \cup \bigcup_{s=i}^{j} Q_s$.

The cycle C satisfies the conditions of Lemma 2, which implies that its edge set is as in (1). Let C' be the unique cycle in $u_i P v_{j-1} \cup \bigcup_{s=i}^{j-1} Q_s$ containing u_i and v_{j-1}. Hence C' is short by the choice of C, and its edge set is given by (1)—with $j-1$ instead of j. Then, $E(C) \Delta E(C')$ is equal to $u_j P v_j \cup Q_j$, which is a short cycle by (E2). As $|E(C)| \leq |E(C')| + |E(C \Delta C')| < 2\ell$, the length of the long cycle C is less than 2ℓ, which contradicts the assumption of the lemma. □

Lemma 4. *Let P be a path in a graph G, and let $(Q_1, \dots, Q_r)$ be a tuple of P-paths that satisfy* (E2)–(E4) *and*

(E1′) *if $|i - j| > 1$, then Q_i and Q_j are internally disjoint.*

If every long cycle in G has length at least 2ℓ, then there is an extension $(Q'_1, \dots, Q'_s)$ of P with $\bigcup_{i=1}^{s} Q'_i \subseteq \bigcup_{j=1}^{r} Q_j$.

Proof. Among all tuples $(Q'_1, \dots, Q'_s)$ of P-paths in $\bigcup_{j=1}^{r} Q_j$ that satisfy (E2)–(E4) and (E1′) choose a tuple $T' = (Q'_1, \dots, Q'_s)$ such that s is minimal. Such a tuple exists as $(Q_1, \dots, Q_r)$ satisfies (E2)–(E4) and (E1′). Let u'_i and v'_i be the endvertices of Q'_i such that $u'_i <_P v'_i$ holds.

Now, assume that there are two paths Q'_i and Q'_j, $j > i$, that share an internal vertex. By (E1′) we have $j = i + 1$. Following Q'_i from u'_i on, let x be the first vertex of $Q'_i - u'_i$ that also belongs to Q'_{i+1}. Now define a new path R as $R = u'_i Q'_i x Q'_{i+1} v'_{i+1}$. The path R is a P-path as x is an internal vertex and its endpoints are u'_i and v'_{i+1}. Furthermore, the length of the cycle $R \cup u'_i P v'_{i+1}$ is at most

$$\ell(Q'_i \cup u'_i P v'_i) + \ell(Q'_{i+1} \cup u'_{i+1} P v'_{i+1}) < 2\ell,$$

which implies that R is short, by assumption.

Now, the tuple $T'' = (Q_1, \dots, Q'_{i-1}, R, Q'_{i+2}, \dots Q'_s)$ satisfies (E2)–(E4) and (E1$'$) as (E2) was just proved, (E3) is trivial, and (E4) and (E1$'$) are inherited from T' as R just combines two consecutive paths of T'. However, T'' uses only $s-1$ paths, which contradicts the choice of T'. Thus, there are no such paths Q'_i, Q'_j that share an internal vertex and hence T' satisfies (E1).

Assume, that T' does not satisfy (E5); that is, there is an i such that $u'_{i+2} <_P v'_i$. By (E4), we have $u'_i <_P u'_{i+1} <_P u'_{i+2}$ and $v'_i <_P v'_{i+1} <_P v'_{i+2}$ which implies

$$u'_i <_P u'_{i+2} <_P v'_i <_P v'_{i+2}.$$

This is the statement of (E4) for the paths Q'_i and Q'_{i+2} which makes Q_{i+1} unnecessary in T. This is again a contradiction to the minimality of s. Thus, the tuple T' satisfies (E1)–(E5) and is therefore an extension of P. $\square$

Lemma 5. *Let P be a path in a graph G, and let $C_1, \dots, C_r$ be a set of short cycles such that*

(i) *$C_i \cap P = u_i P v_i$ for two (not necessarily distinct) vertices u_i, v_i, for $i = 1, \dots, r$;*

(ii) *C_i and C_{i+1} meet outside P for $i = 1, \dots, r-1$; and*

(iii) *$u_i P v_i$ and $u_{i+1} P v_{i+1}$ meet for $i = 1, \dots, r-1$.*

If every long cycle in G has length at least 3ℓ, then there is a short cycle $C \subseteq \bigcup_{i=1}^r C_i$ such that $C \cap P = u_1 P v_r$.

Proof. By induction on r we show that: there is a short cycle $C \subseteq \bigcup_{i=1}^r C_i$ such that $C \cap P = u_1 P v_r$ and such that C contains an edge in $E(C_r) \setminus E(P)$ that is incident with v_r.

The induction starts with $C = C_1$. Now, let C' be such a cycle for $r-1$. For every i, let Q_i be the path $C_i - u_i P v_i$, and let p_i and q_i be its endvertices such that p_i is a neighbour of u_i in C_i and q_i a neighbour of v_i in C_i. We define Q' with endvertices p', q' in the analogous way as $Q' = C' - u_1 P v_{r-1}$.

Assume first that C' and C_r meet outside P. Starting in p' let x be the first vertex in Q' that lies in Q_r. Then put $C = u_1 p' Q' x Q_r q_r v_r \cup u_1 P v_r$ and observe that C satisfies all required properties if, in addition, it is short. This holds, as $\ell(C) \leq \ell(C') + (\ell(Q_r) + \ell(u_r P v_r)) < \ell + \ell = 2\ell$.

Next, assume that Q' and Q_r are disjoint outside P. Since the edge $q' v_{r-1}$ of C' is an edge of C_{r-1} we see that $q' \in V(Q_{r-1})$, which means that Q' and Q_{r-1} have a vertex in common. Starting from p' let y be the first vertex of Q' that lies in Q_{r-1}. Starting from q_r let z be the first vertex in Q_r that lies in Q_{r-1}. Since C_{r-1} and C_r meet outside P, by (ii), there is such a vertex z. Put $C = u_1 p' Q' y Q_{r-1} z Q_r q_r v_r \cup u_1 P v_r$ and observe that, again, C satisfies all required properties if it is short.

We now prove that C is a short cycle. Using (iii), we see that

$$\begin{aligned}\ell(C) &\leq \ell(C') + \ell(C_{r-1}) + \ell(C_r)\\ &< \ell + \ell + \ell = 3\ell,\end{aligned}$$

as C' is short by induction and as the other two terms are smaller than ℓ as well. Thus, the length of the cycle C is smaller than 3ℓ, which means it is a short cycle. $\square$

3.3 Frames

Simonovits' [19] short proof of the Erdős-Pósa theorem rests on a maximal subgraph of the ambient graph G, in which all the disjoint cycles are found. We mimic this approach that also appears in other works [3, 15]. However, in contrast to all such previous approaches, in our case this subgraph is not subcubic, but may have arbitrary high maximum degree.

Any subgraph F of a graph G is a *frame* of G if its minimum degree $\delta(F)$ is at least 2 and if every cycle in F is long. For a frame F of G, we define

- $U(F) = \{v \in V(F) : d_F(v) \geq 3\}$, the set of vertices of degree at least 3 in F; and
- $\mathrm{ds}(F) = \sum_{u \in U(F)} d_F(u)$, the *degree-sum* of F.

In the proof we will choose a frame of maximal degree-sum. The main motivation stems from the fact that large values in $\mathrm{ds}(F)$ yield k edge-disjoint long cycles in F. In the next lemma we collect a number of useful properties about frames.

Lemma 6. *Let F be a frame of maximal degree-sum in a connected graph G. Then*

(i) *F is connected;*

(ii) *if $\mathrm{ds}(F) \geq 84k \log k$ then G contains k edge-disjoint long cycles;*

(iii) *every F-path is short; and*

(iv) *there exists a short path $P = u \dots v \subseteq F$ for every F-path $Q = u \dots v$. This path is unique if every long cycle in G has length at least 2ℓ.*

We need some preparation before we can prove the lemma.

Lemma 7 (Erdős and Pósa [6])**.** *Let G be a multigraph on n vertices with $\delta(G) \geq 3$. Then G contains a cycle of length at most* $\max\{2\log n, 1\}$.

Lemma 8. *Let $k \in \mathbb{N}$ and G be a multigraph with $|E(G)| \geq 42k \log k$ and $\delta(G) \geq 3$. Then G contains k edge-disjoint cycles.*

Proof. We proceed by induction on k. For $k = 1$ the statement holds, since every multigraph with $\delta(G) \geq 3$ contains a cycle.

Let $k \geq 2$. We may assume that $n \geq 2$, as otherwise the statement is trivial. Let C be a shortest cycle in G. Let n_1 and n_2 be the number of vertices with degree 1 and 2 in $G_0 = G - E(C)$, respectively. Thus $n_1 + n_2 \leq \ell(C)$. As long as

G_t contains a vertex of degree 1 or 2, let G_{t+1} arise from G_t by either deleting a vertex of degree 1 or suppressing a vertex of degree 2. Let s be the maximal integer for which G_s is defined. We claim that one of the following statements hold for the transformation from G_t to G_{t+1}.

(i) The number of vertices of degree 1 does not increase and the number of vertices of degree 2 decreases.

(ii) The number of vertices of degree 1 decreases and the number of vertices of degree 2 increases by at most 1.

To see that our claim is true, suppose we deleted a vertex u of degree 1 and let v be the neighbour of u. If $d_{G_t}(v) = 2$, then (i) holds and otherwise (ii) holds. If we suppress a vertex of degree 2, then (i) holds.

It is easy to see that (ii) holds at most n_1 times. Hence (i) holds at most $n_1 + n_2$ times. Observe that $|E(G_t)| = |E(G_{t+1})| - 1$. Therefore, $|E(G_s)| \geq |E(G)| - \ell(C) - 2n_1 - n_2 \geq |E(G)| - 3\ell(C)$.

Let H arise from G_s by deleting isolated vertices. Thus

$$|E(H)| \geq |E(G)| - 3\ell(C). \tag{2}$$

By construction, H does not contain vertices of degree 1 or 2; thus, $\delta(H) \geq 3$ holds or H is empty. We claim that $|E(H)| > 42(k-1)\log(k-1) \geq 0$. If true, H contains in particular an edge, which implies that $\delta(H) \geq 3$. Moreover, we can apply induction to H to find $k-1$ edge-disjoint cycles in H. Since $G - E(C)$ contains a subdivision of H, we therefore obtain together with C in total k edge-disjoint cycles in G.

It remains to prove that $|E(H)| > 42(k-1)\log(k-1)$. We write $m = |E(G)|$ and by $\delta(G) \geq 3$ we have $|V(G)| \leq \frac{2m}{3}$. As C was chosen as the shortest cycle in G, Lemma 7 implies

$$\ell(C) \leq 2\log\left(\frac{2m}{3}\right). \tag{3}$$

Note that the function $x \mapsto x - 6\log\left(\frac{2}{3}x\right)$ is increasing for $x \geq 9$. Since $k \geq 2$, we conclude $\log(28\log k) \leq 6\log k$. Together with $m \geq 42k\log k \geq 9$, we deduce from (2) and (3) that

$$\begin{aligned}
|E(H)| &\geq m - 6\log\left(\frac{2}{3}m\right) \\
&\geq 42k\log k - 6\log\left(28k\log k\right) \\
&\geq 42k\log k - 6\log k - 6\log(28\log k) \\
&\geq 42k\log k - 6\log k - 36\log k \\
&> 42(k-1)\log(k-1).
\end{aligned}$$

This finishes the proof. $\square$

Proof of Lemma 6. For (i), suppose that F has two components A and B. As G is connected, there is an A–B-path P in G that is internally disjoint from F. Thus, $F \cup P$ is a frame, as $F \cup P$ contains the same cycles as F. Since $\mathrm{ds}(F \cup P) > \mathrm{ds}(F)$, we obtain a contradiction to the choice of F.

For (ii), denote by H the multigraph obtained from F by suppressing all vertices of degree 2. Observe that $|E(H)| = \frac{1}{2}\mathrm{ds}(F) \geq 42k\log k$ and $\delta(H) \geq 3$.

Thus, by Lemma 8, H and then also F contains k edge-disjoint cycles. Since all cycles in F are long, the assertion is proved.

For (iii), suppose there is a long F-path Q. Then it can be added to F, since in $F \cup Q$ all cycles are still long. However, $\mathrm{ds}(F \cup Q) > \mathrm{ds}(F)$, which is a contradiction.

For (iv): As F is connected by (i), the distance of u and v in F is finite. If $\mathrm{dist}_F(u, v) \geq \ell$, then any cycle in $F \cup Q$ containing Q is long, which again contradicts (iii) and proves the first part of (iv). If there were two short u–v-paths P_1, P_2 in F, their union $P_1 \cup P_2 \subseteq F$ would contain a cycle of length less than 2ℓ which is short by assumption. This is impossible as F only contains long cycles. □

3.4 Edge-connectivity

The aim of this subsection is to prove Lemma 13 which is an important tool for defining a hitting set in subsection 4.4. Lemmas 9 to 12 only prepare 13.

A well-known result of Mader [12] states that every graph on n vertices with at least $2kn$ edges contains a $(k+1)$-connected subgraph. This is no longer true for (loopless) multigraphs, but holds if we replace connectivity by edge-connectivity.

Lemma 9. *Let $k \in \mathbb{N}$, and let G be a loopless multigraph on n vertices with at least kn edges. Then G contains a $(k+1)$-edge-connected multigraph as a subgraph.*

Proof. We show that

> *every loopless multigraph G on $n \geq 2$ vertices and at least $kn-k+1$ edges contains a $(k+1)$-edge-connected multigraph as a subgraph.* (4)

For $n = 2$, the statement holds, as G is a graph on two vertices with at least $k+1$ edges joining them, which makes G $(k+1)$-edge-connected itself.

For $n \geq 3$, suppose that there is a counterexample to (4). Pick one, H say, with the smallest number n of vertices. As a counterexample, H has a partition $A \cup B$ of its vertex set such that there are at most k edges joining A and B.

Consider first the case, when one of A, B consists of a single vertex, u say. Then $d_H(u) \leq k$. Since $H - u$ is a graph on $n-1 \geq 2$ vertices and has at least $(kn - k + 1) - k = k(n-1) - k + 1$ edges, it follows, by minimality of H, that $H - u$ contains a $(k+1)$-edge-connected subgraph. But such a subgraph is also a subgraph of H, which is impossible.

Thus, both A and B contain at least two vertices. As the graphs $H[A]$ and $H[B]$ do not contain a $(k+1)$-edge-connected subgraph (as H does not), they have at most $k|A| - k$ and $k|B| - k$ edges, by the minimality of H. Then H has at most $k(|A|+|B|) - 2k + k = kn - k$ edges, which is the final contradiction. □

Let G be a multigraph and $k \in \mathbb{N}$. For two vertices $u, v \in V(G)$, we define $u \sim_k v$ if either $u = v$ or if there are k edge-disjoint u–v-paths in G. The transitivity of $\sim_k$ follows from Menger's theorem and thus $\sim_k$ is an equivalence relation.

Lemma 10. *Let G be a multigraph and let A, B be nonempty subsets of distinct equivalence classes of $\sim_k$. Then there is a set X of at most $k-1$ edges separating A and B.*

Proof. Pick $a \in A$ and $b \in B$, and observe that $a \not\sim_k b$. Thus there is an edge set X of size at most $k-1$ that separates a and b in G. Suppose that X fails to separate A from B in G. Then there are $a' \in A$ and $b' \in B$ such that $G-X$ still contains an a'–b'-path. Since X is too small to separate a from a', and b from b', we see that the vertices a, a', b', b belong to the same component in $G-X$, which is a contradiction. □

Before we proceed, let us note that H-paths have the edge version of the Erdős-Pósa property.

Lemma 11 (Mader [13])**.** *Let $k \in \mathbb{N}$, and let H be a submultigraph of a multigraph G. Then there exist either k edge-disjoint H-paths or a set $X \subseteq E(G)$ of size at most $2k-2$ such that $G-X$ does not contain any H-path.*

Lemma 12. *Let $k, p \in \mathbb{N}$, and let $A_1, \ldots, A_p$ be subsets of p distinct equivalence classes of $\sim_k$ in a multigraph G. Then there is an edge set $X \subseteq E(G)$ of size at most $2pk-2$ such that for all distinct $i, j \in \{1, \ldots, p\}$, the multigraph $G-X$ does not contain any A_i–A_j-path.*

Proof. We may assume that $p \geq 2$. Let G' arise from G by identifying for every $i \in \{1, \ldots, p\}$ all vertices in A_i to a single vertex a_i.

Assume first that G' contains a set $X \subseteq E(G')$ such that for all distinct $i, j \in \{1, \ldots, p\}$ the multigraph $G'-X$ does not contain any a_i–a_j-path. Viewing X as a set of edges in G, we observe that $G-X$ does not contain any A_i–A_j-path for any distinct $i, j \in \{1, \ldots, p\}$: indeed, every A_i–A_j-path in $G-X$ corresponds to an a_i–a_j-path in $G'-X$.

Thus, we may assume that any such set X in G' has size strictly larger than $2pk-2$. As a consequence of Lemma 11, there is therefore a set $\mathcal{P}$ of kp edge-disjoint $\{a_1, \ldots, a_p\}$-paths in G'. Define a multigraph $G^{\mathcal{P}}$ on $\{a_1, \ldots, a_p\}$ as vertex set, where a_i and a_j are joined by q edges if $\mathcal{P}$ contains exactly q edge-disjoint a_i–a_j-paths for distinct $i, j \in \{1, \ldots, p\}$. The multigraph $G^{\mathcal{P}}$ has kp edges and is loopless.

Applying Lemma 9, we obtain a k-edge-connected submultigraph of $G^{\mathcal{P}}$. In particular, there are distinct $i, j \in \{1, \ldots, p\}$ such that a_i and a_j are linked by k edge-disjoint paths $Q_1^{\mathcal{P}}, \ldots, Q_k^{\mathcal{P}}$ in $G^{\mathcal{P}}$. Each such path $Q_i^{\mathcal{P}}$ corresponds to a subset of paths in $\mathcal{P}$ whose union contains an a_i–a_j-path Q_i' in G'. Since the paths in $\mathcal{P}$ are pairwise edge-disjoint, this is also the case for the a_i–a_j-paths $Q_1', \ldots, Q_k'$ in G'.

By Lemma 10, there is a set F of at most $k-1$ edges which separate A_i and A_j in G. The set F, seen as edges in G', then separates a_i from a_j, which is impossible because at least one of the paths $Q_1', \ldots, Q_k'$ avoids F. □

Let A be a vertex set in a multigraph G, and let k be a positive integer. An edge set X *k-perfectly separates* A if for every $a, a' \in A$ with $a \not\sim_k a'$ in $G-X$, the vertices a, a' lie in different components of $G-X$. This means, that two vertices either are not connected or there are at least k edge-disjoint paths between them.

Lemma 13. *Let $k \in \mathbb{N}$, and let A be a vertex set in a multigraph G. Then there is a set $X \subseteq E(G)$ of size at most $4(|A|-1)k$ that k-perfectly separates A.*

Proof. We use induction on $|A|$. Let $A_1, \ldots, A_p$ be a partition of A induced by the equivalence classes of $\sim_k$. If $p = 1$, the statement trivially holds as $X = \emptyset$ k-perfectly separates A. In particular, this covers the case $|A| = 1$.

Therefore, we may assume that $p \geq 2$. We apply Lemma 12 to obtain a set $X' \subseteq E(G)$ of size at most $2pk - 2$ that separates A_i from A_j for all distinct i, j. Denote for every $i \in \{1, \ldots, p\}$ by G_i the union of components in $G - X'$ that contain a vertex in A_i, and observe that the G_i are pairwise disjoint by choice of X'. By induction, there is a set $X_i \subseteq E(G_i)$ of size at most $4(|A_i|-1)k$ that k-perfectly separates $A_i \cap V(G_i)$ in G_i. Thus, $X = X' \cup X_1 \cup \ldots \cup X_p$ k-perfectly separates A. Observe that

$$|X| \leq 2pk - 2 + 4(|A| - p)k \leq 4|A|k - 2pk \leq 4(|A| - 1)k,$$

which completes the proof. □

4 Proof of the main theorem

We start with a brief proof sketch. The key trick is to force a gap between short and long cycles: by induction, we can ensure that there are no intermediate cycles, cycles of length between ℓ and 10ℓ. This forces a lot of structure. Repeatedly, we will argue that this or that property is satisfied because otherwise we would find an intermediate cycle.

Throughout we fix a frame F of maximal degree-sum. As every long cycle that is not contained in the frame contains at least one F-path, it is necessary to find structure in the F-paths. To this end, we group F-paths to *hubs*. The hubs together with parts of the frame F form the *hub closures*, which essentially partition the edge set of G. Informally, the hub closures are the largest 2-connected pieces that may contain cycles without also containing a cycle of F.

From the absence of intermediate cycles we will deduce via the path extensions treated in the previous section that no hub closure contains a long cycle. That means that every long cycle in some sense follows along a cycle in F (without actually being contained in F). In particular, it traverses at least two (in fact, at least three) distinct hub closures. To define a candidate hitting set we therefore disconnect hub closures when this is possible with few edges and when this cuts a connection between branch vertices of F. The resulting edge set is either a true hitting set, or we will be able to piece together k edge-disjoint long cycles that all traverse well-connected hub closures in the same way.

Proof of Theorem 1. We define

$$f(k, \ell) = 714k^2 \log k + 10\ell(k-1).$$

We prove by induction on k that

if a graph G does not contain k edge-disjoint long cycles, then it contains an edge set X of size at most $f(k, \ell)$ that meets every long cycle. (5)

Clearly, (5) is true for $k = 1$ as either G contains a long cycle or $X = \emptyset$ meets all long cycles in G. We therefore assume that

$$k \geq 2 \textit{ and that } G \textit{ does not contain } k \textit{ edge-disjoint long cycles.} \tag{6}$$

Suppose G contains a long cycle C of length at most 10ℓ. As $G - E(C)$ contains at most $k-2$ edge-disjoint long cycles, by induction there is a hitting set $X' \subseteq E(G) \setminus E(C)$ for $G - E(C)$ of size at most $714(k-1)^2 \log(k-1) + 10\ell(k-2)$. Observe that $X = E(C) \cup X'$ is a hitting set of G such that

$$\begin{aligned} |X| = |X'| + |E(C)| &\leq 714(k-1)^2 \log(k-1) + 10\ell(k-2) + 10\ell \\ &\leq 714k^2 \log k + 10\ell(k-1) = f(k, \ell). \end{aligned}$$

Thus, we may assume that

$$\textit{every long cycle of } G \textit{ has length more than } 10\ell. \tag{7}$$

We may also assume that every edge of G lies in a long cycle. Otherwise, if $e \in E(G)$ is not contained in any long cycle, then every hitting set of $G - e$ is also a hitting set of G.

Suppose, G is not 2-connected; that is, G contains several blocks. Note that every cycle lies in exactly one block. Since every edge belongs to at least one long cycle, every block contains a long cycle. Let B be a block of G and let k' be the maximal integer such that B contains k' edge-disjoint long cycles. Hence $0 < k' < k-1$, as $G-B$ contains at least one long cycle that is edge-disjoint from every cycle in B. Observe that $G-B$ contains at most $k-k'-1 < k$ edge-disjoint long cycles. We apply our induction hypothesis to B and $G - B$ and obtain a hitting set $X_1 \subseteq E(B)$ in B of size at most $714(k'+1)^2 \log(k'+1) + 10\ell k' \leq 714(k'+1)^2 \log k + 10\ell k'$ and a hitting set $X_2 \subseteq E(G) \setminus E(B)$ of size at most $714(k-k')^2 \log k + 10\ell(k-k'-1)$. Trivially $X = X_1 \cup X_2$ is a hitting set in G such that

$$\begin{aligned} |X| &\leq 714(k'+1)^2 \log k + 10\ell k' + 714(k-k')^2 \log k + 10\ell(k-k'-1) \\ &\leq 714 \log k \left(k'^2 + 2k' + 1 + k^2 - 2kk' + k'^2\right) + 10\ell(k-1) \\ &= 714 \log k \left(2k'(k'+1-k) + 1 + k^2\right) + 10\ell(k-1) \\ &\leq 714k^2 \log k + 10\ell(k-1) = f(k, \ell) \end{aligned}$$

as $2k'(k'+1-k) + 1 \leq 0$ holds because of $k' < k-1$. Thus, we can assume that

$$G \textit{ is 2-connected.} \tag{8}$$

We now choose a frame F of G of maximal degree-sum $\mathrm{ds}(F)$ (and we may assume that G contains at least one long cycle, which implies that a frame in G exists), which we let be fixed throughout the whole proof. As F only contains long cycles, (7) implies that

$$\textit{the girth of } F \textit{ is more than } 10\ell. \tag{9}$$

Next, we investigate $G - F$ and how the components of $G - F$ attach to F.

4.1 Bridges of the frame

In the light of (6) and (7), Lemma 6 now states:

$$F \text{ is connected; } \mathrm{ds}(F) < 84k \log k\text{; every } F\text{-path } Q = u \dots v \text{ is short and } F \text{ contains a unique short } u\text{–}v\text{-path } P. \tag{10}$$

For any F-path $Q = u \dots v$, we call the unique short u–v-path in F its *shadow* and denote it by S_Q.

An *F-bridge* of G or simply a *bridge* is either an edge in $E(G) \setminus E(F)$ with its two endvertices in $V(F)$, or a component K of $G - F$ together with all its neighbours N in F and all edges of G joining K and N. Equivalently, a bridge is the union of all F-paths that form a component in the graph on the set of all F-paths where two F-paths are adjacent if the share an internal vertex. For an F-bridge B of G, we call the vertices in $B \cap F$ the *feet* of B (in H). The *shadow* S_B of B is the union of the shadows of all F-paths contained in B.

Claim 1. *For every bridge B, the shadow S_B is a tree of diameter less than ℓ.*

Proof. As B is connected, it contains an x–y-path Q between any two of its feet x, y. The shadow of this F-path Q connects x and y in S_B. As all vertices in S_B that are no feet lie in the shadow of an F-path between two feet, we conclude that S_B is connected.

Suppose that S_B contains a cycle C. Since C is contained in F, it follows that C is a long cycle, which, in turn, implies $\ell(C) \geq 10\ell$, by (7). Pick two vertices r_1, r_2 in C at distance precisely 2ℓ in C, and let R be the subpath of C of length 2ℓ between r_1 and r_2.

Why is r_i in S_B? Because there is an F-path $Q_i \subseteq B$ whose shadow P_i contains r_i. Denote by x_i, y_i the endvertices of P_i, and observe that P_i is a short path, by Lemma 6 (iv). By the same statement, there exists also a short x_1–x_2-path S in the shadow of B.

Since $P_1 \cup P_2 \cup R \cup S \subseteq S_B$ has at most 5ℓ edges it cannot contain a long cycle, and because it is a subset of F it cannot contain a short cycle. In particular, this means that $S = x_1 P_1 r_1 R r_2 P_2 x_2$, and thus that $R \subseteq S$. This, however, is impossible since S has length at most ℓ but R has length 2ℓ.

We deduce that S_B is a tree. By the definition of a shadow, every leaf of S_B is a foot. As any two feet of B are connected by an F-path, their distance in S_B is short by Lemma 6 (iv). Thus, the diameter of S_H is less than ℓ. □

4.2 Hubs

We define a graph $\mathcal{G}$ on the set of all bridges of G, where two bridges B_1, B_2 are adjacent if their shadows share a common edge. A *hub* is the union of all bridges in a component in $\mathcal{G}$. Thus, a hub is a subgraph of G consisting of all bridges that form a component in $\mathcal{G}$. We say that a bridge B *belongs* to a hub H if $B \subseteq H$, that is, if B is part of the component in $\mathcal{G}$ that defines H. For a hub H, the *shadow* S_H of H is the union of the shadows of all bridges in H. By Claim 1, the graph S_H is connected. We will write $\overline{H}$ for $H \cup S_H$ and call it the *closure of H*.

One key step in our main proof is Claim 8 where we show that a hub closure does not contain a long cycle. To this end, we first show that the shadow of a hub does not contain a (long) cycle.

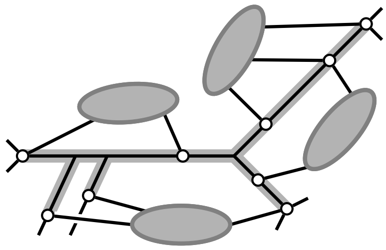

Figure 3: A hub consisting of four bridges, and its shadow (in grey).

Let us start with a simple observation.

Claim 2. *For every hub H, the closure $\overline{H}$ is 2-connected.*

Proof. Since G is 2-connected, a bridge together with its shadow is 2-connected, too. The closure of a hub is the union of adjacent bridges together with their shadows. As adjacent bridges overlap on an edge, the union again is 2-connected. □

For a hub H, let L_H be the graph with vertex set $E(S_H)$ and $e, f \in V(L_H)$ are adjacent in L_H if e, f share a common vertex in G and there is a bridge B which belong to H such that $e, f \in E(S_B)$. Let L_H^* arise from L_H by adding all possible edges of the following type: for all $e_1, \ldots, e_r \in V(L_H)$ sharing a common vertex in G which induce a connected graph in L_H add all edges e_ie_j for $i, j \in \{1, \ldots, r\}$.

Claim 3. *The graph L_H^* is connected for every hub H.*

Proof. We will prove that L_H is connected which immediately proves the claim as $L_H \subseteq L_H^*$. First, it is easy to see that for any bridge B of H, the induced subgraph $L_H[E(S_B)]$ on the edges of S_B is connected. This holds as edges of S_B with common endvertex in G are adjacent in L_H as they belong to the shadow of the same bridge. The connectivity of S_B then implies the connectivity of $L_H[E(S_B)]$.

Let $e, f \in V(L_H)$ be two edges of the hub H that belong to the shadows of different bridges B, B'. The definition of hubs implies that there is a sequence of bridges $B = B_1, B_2, \ldots, B_r = B'$ such that S_{B_i} and $S_{B_{i+1}}$ share at least one edge. As all $L_H[E(S_{B_i})]$ are connected in L_H, there is a path in L_H joining e and f. □

Let $\mathcal{E} = (Q_1, \ldots, Q_r)$ be an extension of a path P. To simplify notation we will identify the graph $\bigcup_{i=1}^r Q_i \cup P$ with the tuple $\mathcal{E} = (Q_1, \ldots, Q_r)$ (bending the definition a bit). Thus, it will make sense to speak of vertices in an extension. The following two claims are a bit technical but provide tools to prove that shadows and closures of hubs do not contain long cycles.

Claim 4. *Let H be a hub, and let P be a path in S_H such that every P-path in F has length at least 3ℓ and such that every pair of consecutive edges in P is adjacent in L_H^*. Then there is an extension $\mathcal{E}$ of P that is contained in $\overline{H}$ and for which $\mathrm{dist}_F(u, P) \leq \ell$ holds for every $u \in V(\mathcal{E}) \cap V(F)$.*

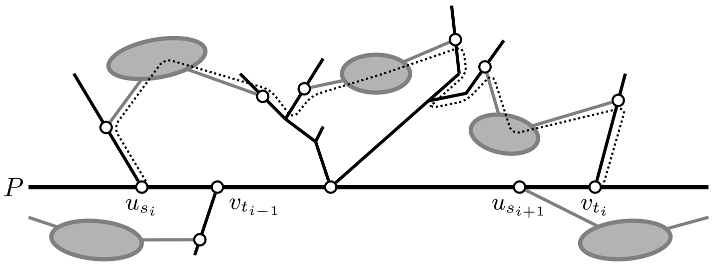

Figure 4: The path Q_i (dotted).

Proof. As before, denote by $\leq_P$ the order on the vertices of P induced by the path, where we fix arbitrarily one of the two endvertices as first vertex.

Denote by P' the union of $E(P)$ and all edges in F that have an endvertex in P. By assumption, the set P' (seen as a vertex set in L_H) contains a path in L_H that contains $E(P)$ entirely (recall that two consecutive edges of P may be nonadjacent in L_H, but adjacent in L^*_H). That means, there is a sequence of bridges $B_1, \ldots, B_t$ such that

$$P \subseteq \bigcup_{i=1}^{t} S_{B_i}, \textit{ and } E(S_{B_i} \cap S_{B_{i+1}}) \cap P' \neq \emptyset \textit{ for } i = 1, \ldots, t-1. \tag{11}$$

We choose the sequence $B_1, \ldots, B_t$ such that t is minimal. Moreover, we fix that the shadow of the first bridge B_1 contains the first edge of P (and then the shadow of B_t contains the last edge of P). To avoid double subscripts we write S_i for the shadow S_{B_i}.

We quickly note:

for every bridge B, the intersection $S_B \cap P$ is a subpath of P. (12)

Indeed, this is the case as S_B is connected and of diameter less than ℓ (Claim 1) and as there are no P-paths in F of length at most 3ℓ, by assumption.

We need a claim about the start and end of P:

if S_i contains the first vertex of P, then $i = 1$, and if S_i contains the last vertex of P, then $i = t$. (13)

Suppose that S_i contains the first vertex of P and that $i > 1$. Then, omitting the bridges $B_1, \ldots, B_{i-1}$ we still have a sequence of bridges that satisfies (11); that P is still contained in the union of the shadows is due to (12). But this contradicts the minimal choice of $B_1, \ldots, B_t$. The argument for the last vertex of P is symmetric.

We claim:

if $|i - j| > 1$, then $S_i \cap S_j$ is either empty or consists of a single vertex in P. (14)

Let $S_i \cap S_j$ be non-empty and $i < j - 1$. Suppose first that S_i and S_j contain a common edge e that lies in P'. Then we could omit the bridges $B_{i+1}, \ldots, B_{j-1}$ from the sequence and still retain (11); that P is still contained in the union of the shadows is due to (12).

Next, suppose that $S_i \cap S_j$ contains a vertex v outside P. Both shadows, which are contained in F, contain a v–P-path of length at most ℓ, by Claim 1.

As we had assumed that there are no P-paths in F of length at most 3ℓ, this implies that $S_i \cap S_j$ contains a v–P-path, which in turn means that $S_i \cap S_j$ contains an edge in P', which is impossible as we have seen. Thus, $S_i \cap S_j \subseteq P$.

By (12), the set $S_i \cap S_j = S_i \cap S_j \cap P$ is a subpath of P. If it contains more than one vertex, it thus contains an edge in P', which we had already excluded. This proves (14).

For every $i = 1, \dots, t-1$ pick an edge e_i in $S_i \cap S_{i+1} \cap P'$—this is possible, by (11). Denote by e_0 the first edge of P, and by e_t the last edge of P. For every $i = 1, \dots, t$, there is, by Claim 1, a path in S_i containing e_{i-1} and e_i. Let S'_i be a longest such path. By definition of a shadow, the endvertices of S'_i are feet of B_i. Pick a path through B_i and use it to complete S'_i to a cycle C_i.

We claim:

(i) *$C_i \subseteq S_i \cup B_i$ is a short cycle;*

(ii) *there are vertices $u_i \leq_P v_i$ such that $u_i P v_i = C_i \cap P$;* (15)

(iii) *C_i and C_{i+1} meet in an edge of P'; and*

(iv) *$u_i P v_j \subseteq \bigcup_{s=i}^{j} S_s$ for every $1 \leq i \leq j \leq t$.*

That C_i is short follows from (10), Claim 1 and (7); (ii) follows from (12), and (iii) holds since both C_i and C_{i+1} contain the edge e_i. Finally, (iv) is a consequence of (i), (ii) and (iii).

We also note that since $e_0 \in E(C_1)$ and $e_t \in E(C_t)$:

$$u_1 \text{ is the first vertex of } P\text{, and } v_t \text{ is its last.} \tag{16}$$

The intersections of $C_i \cap P = u_i P v_i$ are paths. Two such paths of consecutive cycles C_i and C_{i+1} may intersect in a single vertex or in a longer path (they meet by (15) (iii)). Let $s_2 < \dots < s_r$ be precisely those indices such that $C_{s_i - 1} \cap C_{s_i} \cap P$ contains at least one edge. For a slightly less cumbersome notation, define also $t_{i-1} = s_i - 1$ and set $s_1 = 1$ and $t_r = t$. Then the cycles $C_1, \dots, C_t$ partition into sets $\{C_{s_i}, \dots, C_{t_i}\}$ for $i = 1, \dots, r$ such that always $C_{t_{i-1}}$ and C_{s_i} share an edge of P. We claim:

for $i = 1, \dots, r$, there is a P-path $Q_i \subseteq \bigcup_{s=s_i}^{t_i} (B_s \cup S_s)$ between u_{s_i} and v_{t_i} such that $Q_i \cup u_{s_i} P v_{t_i}$ is a short cycle. (17)

We prove this with Lemma 5 and therefore check that the conditions of Lemma 5 are satisfied. The first condition follows from (ii). Why do C_s and C_{s+1} for $s \in \{s_i, \dots, t_i - 1\}$ meet outside P? Because C_s and C_{s+1} have a common edge e in P' by (15) that, however, e cannot lie in P by definition of the s_i. Thus, the endvertex of e outside P is a common vertex that lies outside P. The other endvertex of e, the one in P, shows that C_s and C_{s+1} meet also in P. Now, the application of the lemma yields a short cycle $C \subseteq \bigcup_{s=s_i}^{t_i} (B_s \cup S_s)$ such that $C \cap P = u_{s_i} P v_{t_i}$. As S_{s_i} needs to contain an edge of P, by definition of s_i, we deduce that $u_{s_i} <_P v_{t_i}$, and in particular that $u_{s_i} \neq v_{t_i}$. Deleting all vertices of C in the interior of $u_{s_i} P v_{t_i}$ results in the desired P-path Q_i.

We note rightaway:

every vertex of F in $\bigcup_{i=1}^{r} Q_i$ has distance at most ℓ from P in F. (18)

Indeed, such a vertex in F lies in some shadow S_s. Every such shadow meets P, by (11), and has diameter at most ℓ (Claim 1), which results in a distance at most ℓ to P in F since $S_s \subseteq F$.

Next:

$$\textit{if } |i-j| > 1, \textit{ then } Q_i \textit{ and } Q_j \textit{ are internally disjoint.} \tag{19}$$

Since two distinct bridges that meet meet in their shadows, we obtain that $Q_i \cap Q_j$ is contained in

$$\left(\bigcup_{s=s_i}^{t_i} S_s\right) \cap \left(\bigcup_{s=s_j}^{t_j} S_s\right),$$

which is contained in P by (14) as $|s_j - t_i| > 1$ since $|j-i| > 1$. Since Q_i and Q_j are P-paths they can thus only meet in their endvertices. This proves (19).

Next:

$$u_{s_i} <_P u_{s_{i+1}} <_P v_{t_i} <_P v_{t_{i+1}} \textit{ for } i = 1, \ldots, r-1. \tag{20}$$

We prove this by induction on i. By definition of the s_i, the paths $u_{s_i}Pv_{t_i}$ and $u_{s_{i+1}}Pv_{t_{i+1}}$ have a common edge. This implies $u_{s_i} <_P v_{t_{i+1}}$.

Suppose that $u_{s_{i+1}} \leq_P u_{s_i}$. Then $i > 1$, by (13) and (16). By induction, we get $u_{s_{i-1}} <_P u_{s_i} <_P v_{t_{i-1}}$. Since we also have that $u_{s_{i+1}} \leq_P u_{s_i} <_P v_{t_{i+1}}$, we deduce that $u_{s_{i-1}}Pv_{t_{i-1}}$ and $u_{s_{i+1}}Pv_{t_{i+1}}$ have a common edge. By (15) (iv), this means that there are $s \in \{s_{i-1}, \ldots, t_{i-1}\}$ and $s' \in \{s_{i+1}, \ldots t_{i+1}\}$ such that S_s and $S_{s'}$ have an edge in common—but this contradicts (14). Thus, we get

$$u_{s_i} <_P u_{s_{i+1}} <_P v_{t_i},$$

because $u_{s_i}Pv_{t_i}$ and $u_{s_{i+1}}Pv_{t_{i+1}}$ have a common edge. Suppose that $v_{t_{i+1}} \leq_P v_{t_i}$. By (13) and (16), this implies $t_{i+1} < t$, which in turn implies $i+1 < r$. Moreover, $u_{s_{i+1}}Pv_{t_{i+1}} \subseteq u_{s_i}Pv_{t_i}$. By definition of the s_i, it follows that $u_{s_i}Pv_{t_i}$ and $u_{s_{i+2}}Pv_{t_{i+2}}$ have an edge in common. Again from (15) (iv) we get that there is an $s \in \{s_i, \ldots, t_i\}$ and an $s' \in \{s_{i+2}, \ldots t_{i+2}\}$ such that u_sPv_s and $u_{s'}Pv_{s'}$ share an edge. Since this edge then lies in the shadow S_s and in the shadow $S_{s'}$, we obtain again a contradiction to (14). This proves (20).

We now apply Lemma 4 to $Q_1, \ldots, Q_r$ in order to obtain the desired extension of P. We note that (16), (17), (19) and (20) ensure that all conditions are satisfied. The desired extension of P does not contain vertices v of F such that $\mathrm{dist}_F(u, P) > \ell$ because of (18). □

Claim 5. *Let H be a hub. Let $P \subseteq F$ be a path of length at most 5ℓ with first and last edge e and f such that e and f belong to S_H but are not adjacent in L_H^*. Then there is no e–f-path in $L_H^* - (E(P) \setminus \{e, f\})$.*

Proof. Suppose there is such a $P \subseteq F$ and an e–f-path Q^* in $L_H^* - (E(P) \setminus \{e, f\})$. Among all such pairs (P, Q^*) choose P and Q^* such that $\ell(Q^*)$ is minimal. We claim that

$$e \textit{ and } f \textit{ are not adjacent in } G. \tag{21}$$

If $\ell(P) = 5\ell$ then obviously e and f cannot be adjacent. Suppose that $\ell(P) < 5\ell$ and let e' be the successor of e in Q^* (note that e' is a vertex in Q^* but an edge in G). We construct a path P' that contradicts the minimal choice of P together with $Q^* - e$. As $\ell(P) < 5\ell$, the graph $P + e'$ cannot contain a cycle because of $P + e' \subseteq F$ and (9).

If $P + e'$ is a path, set $P' = P + e'$. Since e' and f are not adjacent in G they are not in L^*_H either. If $P + e'$ is not a path, set $P' = P - e + e'$. If e' and f were adjacent in L^*_H, either $P + e'$ contained a cycle or $P = ef$ and e, f, e' all share a common vertex—then, however, the definition of L^*_H implies that e and f have to be adjacent in L^*_H, too, which we have excluded.

As $\ell(P) < 5\ell$, the new path P' satisfies $\ell(P') \leq 5\ell$ and there is a path in $L^*_H - (E(P') \setminus \{e', f\})$ joining its endvertices e' and f, namely $Q^* - e$. Thus, $(P', Q^* - e)$ contradicts the minimality of Q^*. This proves (21).

Consider the subgraph Q of G that consists of the edges $V(Q^*)$ and all incident vertices. We claim that Q is a path. By the definition of L^*_H, Q is connected. Thus, if Q is not a path, it contains a vertex v of degree at least 3. Starting with e, let e' be the first vertex of Q^* that, seen as an edge in G, contains v as an endvertex and let f' be the last such vertex of Q^*. As $d_Q(v) \geq 3$, the edges e' and f' are not adjacent in L^*_H as Q^* was chosen minimal. Note that the path $e'Q^*f'$ in L^*_H is shorter than Q^* as $\{e', f'\} \neq \{e, f\}$, by (21). Thus, the path $P' = e'f'$ together with the path $e'Q^*f'$ in L^*_H form a pair $(P', e'Q^*f')$ that contradicts the minimality of Q^*. Therefore, Q is a path in G.

Our next aim is to find a subpath $Q' \subseteq Q$ that satisfies the following two conditions:

$$\textit{every } Q'\textit{-path in } F \textit{ has length at least } 5\ell\textit{; and} \tag{22}$$

$$\textit{there is a } Q'\textit{-path } R \subseteq F \textit{ between the endvertices of } Q' \textit{ of length } 5\ell. \tag{23}$$

The set of those subpaths that satisfy (22) is nonempty, since every subpath of Q of length, say, at most ℓ satisfies (22)—recall that the girth of F is larger than 10ℓ by (9).

Pick a longest subpath S of Q that satisfies (22) in the role of $S = Q'$. If S also satisfies (23), we found the desired path. Thus, we may assume that the shortest S-path $R \subseteq F$ between the endvertices u and v of S has length larger than 5ℓ. Suppose that u, v are precisely the endvertices of Q. Since $\ell(P) \leq 5\ell$, either P is a shorter S-path than R, which is impossible, or P contains a S-path of length less than 5ℓ, which violates (22). Therefore, at least one of u, v is not an endvertex of Q; let this be u.

Thus, S can be extended by the unique neighbour u' of u in $V(Q) \setminus V(S)$ to a path $S' \subseteq Q$. By the maximality of $\ell(S)$, the path S' does not satisfy (22). This is only possible if there is an S'-path $R' \subseteq F$ between u' and some vertex $y \in V(S)$ that has length less than 5ℓ. Since $R = uu'Ry$ is an S-path in F it follows from (22) that

$$5\ell \geq \ell(R) = \ell(R') + 1 \geq 5\ell - 1 + 1 = 5\ell.$$

Setting $Q' = uSy$ yields a subpath of Q satisfying (22) and (23).

Let x and y be the endvertices of Q' (and of R). We check that the conditions of Claim 4 are satisfied by Q'. As Q' satisfies (22), every Q'-path in F has length at least 3ℓ. The path Q' is a subpath of Q for which $E(Q)$ is a path in L^*_H, and thus Q^* is also a path in L_H^*. This implies the second condition of the claim. Thus, by Claim 4, there is an extension $\mathcal{E} \subseteq \overline{H}$ of Q' that uses no vertex of F at distance more than ℓ from Q' measured in F. By Lemma 3, G either contains a long cycle of length at most 2ℓ, which is impossible by (7), or every cycle in $\mathcal{E}$ is short. Thus the cycle $C \subseteq \mathcal{E}$ containing x and y is short; recall that Lemma 2 ensures that there is such a cycle C.

Denote by R' the path obtained from R by removing the first $\ell + 1$ vertices and the last $\ell + 1$ vertices. Note that $\ell(R') \geq 3\ell - 2 \geq 2\ell$ as $\ell(R) = 5\ell$. We claim that every vertex of R' has distance more than ℓ from Q' in F. Suppose not. Then there exists a Q'–R'-path P_1 of length at most ℓ. From the endvertex of P_1 in R' pick a subpath P_2 of R that ends in x or in y and has length at most 3ℓ which is possible as $\ell(R) = 5\ell$. Since $P_1 \neq P_2$ the union $P_1 \cup P_2$ either contains a cycle or is a Q'-path. It cannot contain a cycle, since such a cycle would be contained in F but would have length at most $\ell(P_1) + \ell(P_2) \leq 4\ell$. Thus, $P_1 \cup P_2 \subseteq F$ is a Q'-path of length at most 4ℓ—this contradicts (22) and hence $\mathrm{dist}_F(R', Q') > \ell$.

Since the cycle C does not contain any vertex in F at distance more than ℓ measured in F, it follows that C is disjoint from R'. We extend R' to a subpath R'' of R that is a C-path. Then, R'' has length

$$2\ell \leq \ell(R') \leq \ell(R'') \leq \ell(R) = 5\ell.$$

Consequently, as C is short, each of the two cycles in $C \cup R''$ through R'' then have length between 2ℓ and $\ell(C) + \ell(R) \leq \ell + 5\ell = 6\ell$, which is impossible by (7). Thus, there are no counterexamples to the claim. □

Using the previous claim, we show that some assumptions of Claim 4 are always satisfied and thus we obtain a simpler version of Claim 4.

Claim 6. *Let H be a hub. Then*

(i) *every pair of edges $e, f \in E(S_H)$ with a common endvertex is adjacent in L_H^*;*

(ii) *every S_H-path in F has length at least 4ℓ; and*

(iii) *for every path $P \subseteq S_H$, there is an extension $\mathcal{E}$ of P that is contained in $\overline{H}$ such that $\mathrm{dist}_F(u, P) \leq \ell$ holds for every $u \in V(\mathcal{E}) \cap V(F)$.*

Proof. For a proof of (i), let $e = uv, f = vw \in E(S_H)$ share an endvertex but be non-adjacent in L_H^*. Then $P = uvw$ is a path in F of length $2 \leq 5\ell$. Applying Claim 5 to P, we see that there is no e–f-path in L_H^*, which is impossible as L_H^* is connected, by Claim 3.

To see (ii), suppose there is an S_H-path $P = u \dots v$ in F of length less than 4ℓ. Let $e, f \in E(S_H)$ be such that e and f contain u and v, respectively. The edges e and f cannot share an endvertex because then F would contain a cycle $P + e + f$ of length less than 5ℓ which contradicts (9). In particular, e and f are not adjacent in L_H^*. Extend P by these two edges and apply Claim 5 in order to obtain a contradiction to L_H^* being connected (Claim 3).

Statement (iii) is exactly the statement of Claim 4 without the assumptions that P induces a path in L_H^* (which is satisfied by (i)) and that P-paths in F have length at least 4ℓ (which is satisfied by (ii)). □

Claim 7. *For every hub H, the graph S_H is a tree.*

Proof. We observed before that S_H is connected. Since $S_H \subseteq F$, it follows that S_H is acyclic unless it contains a long cycle C.

By Claim 6 (i), every two consecutive edges in C are adjacent in L_H^* which means that $E(C)$ is a cycle in L_H^*. Then, however, we obtain a contradiction to Claim 5 with any path $P \subseteq C$ of length 3. □

Claim 8. *For every hub H, the closure $\overline{H}$ of H does not contain a long cycle. Thus, the diameter of $\overline{H}$ is at most $\frac{\ell}{2}$.*

Proof. Suppose there is a long cycle C in $\overline{H}$. We say that C *traverses* a F-bridge B of H t times if $C \cap B$ contains exactly t non-trivial components (where non-trivial means that the component contains an edge). Among all long cycles in $\overline{H}$ choose C such that the total number of bridge traversals of C is minimal. We will prove that C does not traverse any bridge.

Suppose that C traverses a bridge B. Let $P = u \dots v \subseteq C \cap B$ be a non-trivial F-path. Assume first that the intersection $C \cap S_P$ of C and the shadow of P contains only one component. As $u, v \in V(C \cap S_P)$, this implies that $S_P \subseteq C$ and together with $P \subseteq C$, we have $C = P \cup S_P$. Then C would have length at most 2ℓ, since F-paths as well as their shadows are short by Lemma 6. This, together with (7), contradicts the assumption that C is long.

Hence we may assume that $C \cap S_P$ contains at least two components. Let $Q \subseteq S_P$ be a C-path in S_P that joins two components of $C \cap S_P$. In $C \cup Q$, there are two cycles D_1 and D_2 that both contain Q. Let D_1 be the one that contains P. As $\ell(C) \geq 10\ell$ by (7), one of the two cycles D_1 or D_2 has length at least 5ℓ and is thus long.

If D_2 is long, it contradicts the choice of C as D_2 traverses B fewer times than C and no other bridge more often than C. Otherwise D_2 is short and as F does not contain short cycles, D_2 traverses a bridge B' (that is not necessarily distinct from B). Then, the other cycle D_1 traverses B' fewer times than C and no other bridge more often than C. As D_1 is long when D_2 is short, the cycle D_1 contradicts the choice of C.

This implies that C does not traverse any bridge of H; that is, $C \subseteq S_H$. This, however, is a contradiction to Claim 7 and we conclude that there is no long cycle in $\overline{H}$.

For any distinct $u, v \in V(\overline{H})$, there is a cycle $C \subseteq \overline{H}$ through u and v, as $\overline{H}$ is 2-connected by Claim 2. Since hub closures do not contain long cycles, the cycle C has length at most $\ell - 1$. Thus, $\mathrm{dist}_{\overline{H}}(u, v) \leq \mathrm{dist}_C(u, v) \leq \frac{\ell}{2}$. This implies $\mathrm{diam}(\overline{H}) \leq \frac{\ell}{2}$. □

Claim 9. *Let H be a hub, and let $u, v \in V(S_H)$, $u \neq v$. Let $r \in \mathbb{N} \cup \{0\}$ be such that the unique u–v-path $Q \subseteq S_H$ has length*

$$r\ell < \ell(Q) \leq (r+1)\ell.$$

Then, for any $t \in \{0, \dots, r\}$, there is a u–v-path $P \subseteq \overline{H}$ such that

$$t\ell \leq \ell(P) \leq t\ell + \frac{3}{2}\ell$$

and $\mathrm{dist}_F(w, Q) \leq \ell$ for every vertex $w \in V(F \cap P)$.

Proof. If $t = r$, then we can choose $P = Q$. Suppose therefore that $t \in \{0, \dots, r-1\}$. Note that $\ell(Q) > r\ell \geq (t+1)\ell$. Let $x \in V(Q)$ be the vertex on Q with $\ell(uQx) = (t+1)\ell$. Next we use Claim 6 (iii) to obtain an extension $\mathcal{E} \subseteq \overline{H}$ of xQv that uses no vertices of F at distance more than ℓ from xQv measured in F. By Lemma 2, the extension $\mathcal{E}$ contains a cycle C through x and v. Lemma 3 together with (7) implies that C is short. Thus, there exists a x–v-path $R \subseteq C$ in $\mathcal{E}$ of length at most $\frac{\ell}{2}$.

Starting from u, let y be the first vertex of Q that lies in R. Note that $t\ell \leq \ell(uQy) \leq \ell(uQx) = (t+1)\ell$, as $\mathrm{dist}_F(y, xQv) \leq \ell$, because $y \in V(\mathcal{E})$. Thus the path $P = uQyRv$ is a path in $\overline{H}$ such that

$$t\ell \leq \ell(uQy) \leq \ell(P) \leq \ell(uQx) + \ell(R) \leq (t+1)\ell + \frac{\ell}{2}.$$

Every vertex w of P has either distance 0 from Q (if $w \in V(uQy)$) or at most ℓ (if $w \in V(yRv)$ as $R \subseteq \mathcal{E}$) in F. □

4.3 Gates of hubs

For a hub H, we call those vertices $v \in V(S_H)$ that have neighbours in $F - S_H$ the *gates* of H. Equivalently, v is a gate of H if it lies in $\overline{H}$ and has a neighbour outside $\overline{H}$. Thus, every path in G that contains a vertex in $G - \overline{H}$ and a vertex in $\overline{H}$ also contains a gate of H.

Claim 10. *The shadows S_{H_1} and S_{H_2} of any two distinct hubs H_1, H_2 share at most one vertex.*

Proof. Since two bridges belong to the same hub if their shadows share an edge, it follows that $S_{H_1} \cap S_{H_2}$ is a collection of isolated vertices. In particular, every vertex of $S_{H_1} \cap S_{H_2}$ is a gate of H_1 and of H_2.

Suppose that $|V(S_{H_1} \cap S_{H_2})| \geq 2$. Consider two common gates g, g' of H_1 and H_2. Pick a g–g'-path P in F that is shortest among all paths contained in F. Then P either contains a S_{H_1}-path or a S_{H_2}-path (or both), which then, by Claim 6 (ii), has length at least 4ℓ. Thus

every two common gates g, g' of H_1 and H_2 have distance at least 4ℓ in F and therefore also in S_{H_i}. (24)

Among all paths that join two common gates of H_1 and H_2, let R be the shortest such path, and let g, g' be its endpoints. Observe that every common gate h distinct from g, g' of H_1 and H_2 is at distance at least 2ℓ from R in S_{H_1}; otherwise, by (24), there would be a shorter path between two common gates. Observe that also in F the common gate h has distance at least 2ℓ from R: otherwise, F would contain an S_{H_1}-path of length at most 2ℓ, contradicting Claim 6 (ii).

Extend R with a g–g'-path through S_{H_2} to a cycle C. As both g–g'-paths in C have length at least 4ℓ, we apply Claim 9 with $t = 1$ and obtain g–g'-paths $P_1 \subseteq \overline{H}_1$ and $P_2 \subseteq \overline{H}_2$ each of length at least ℓ and at most 3ℓ. In addition, the claim ensures that every vertex in $P_1 \cap F$ has distance at most ℓ measured in F to R. As $P_i \subseteq \overline{H_i}$ for $i = 1, 2$, every vertex of $V(P_1) \cap V(P_2)$ is a common gate of H_1 and H_2, which then has distance at most ℓ to C in F. Any common gate other than g or g' has distance at least 2ℓ to R in F, as argued above. Thus,

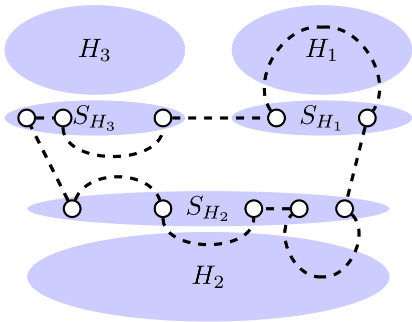

Figure 5: The dashed cycle traverses and visits H_1 once, it traverses H_2 once and visits H_2 twice. It does not traverse H_3, thus it also does not visit it.

P_1 and P_2 meet only in g, g' and $P_1 \cup P_2$ is a cycle. The length of the cycle is between 2ℓ and 6ℓ, which is impossible by (7). Therefore, S_{H_1} and S_{H_2} meet in at most one gate. $\square$

Claim 11. *Let C be a cycle, and let H be a hub such that C contains an edge both in $E(\overline{H})$ and in $E(G) \setminus E(\overline{H})$. Then C is long.*

Proof. We say a cycle C *traverses* a hub H if C contains an edge of H. The number of traversals of H is the number of components of $C \cap \overline{H}$ that contain an edge of H. For hubs H that are traversed by C, we define the number of *visits* as the number of components of $C \cap \overline{H}$ (which will be larger than the number of traversals if $C \cap \overline{H}$ has components that are contained in the shadow of H). When C fails to traverse H then the number of visits is 0.

Suppose there is a short cycle that contains an edge of some hub closure $\overline{H}$ but is not completely contained in $\overline{H}$. Choose such a cycle C such that the total number of hub traversals is minimal and subject to that choose C such that the total number of visits is minimal.

We claim:

$$\textit{if } C \textit{ traverses a hub } H\textit{, then } C \cap \overline{H} \textit{ is a path.} \tag{25}$$

Suppose that $C \cap \overline{H}$ has a component Q_1 with an edge in H (as C traverses H) and a second component (with or without edge in H). By Claim 8, the diameter of $\overline{H}$ is at most $\frac{\ell}{2}$. Thus, there is a C-path $P \subseteq \overline{H}$ of length at most $\frac{\ell}{2}$ that starts in Q_1 and ends in another component Q_2 of $C \cap \overline{H}$. Let D_1, D_2 be the two cycles in $C \cup P$ that contain P. We observe that

$$\textit{each of } D_1, D_2 \textit{ shares an edge with } \overline{H} \textit{ but is not contained in } \overline{H}. \tag{26}$$

Indeed, each of D_1, D_2 shares an edge with $\overline{H}$ because of $P \subseteq \overline{H}$. Neither of D_1, D_2 is contained in $\overline{H}$: running along the P-path $D_i \cap C$ from the endvertex of P in Q_1 we see that the first edge outside Q_1 lies also outside $\overline{H}$, and there must be such an edge since Q_1 and Q_2 are distinct components.

Moreover,

$$\textit{each of } D_1, D_2 \textit{ is a short cycle.} \tag{27}$$

For $i = 1, 2$, the length of D_i is at most $\ell(C) + \ell(P) \leq \ell + \frac{\ell}{2}$. As every long cycle has length at least 10ℓ by (7), we deduce that D_i is short.

The cycles D_1, D_2 are thus also counterexamples of the claim. To see that one of them contradicts the minimal choice of C, we distinguish two cases.

First, assume that C traverses a second hub $H' \neq H$. Then one of D_1, D_2, say D_1, meets an edge of H'. It follows that D_2 has at least one hub traversal less than C and, in light of (26) and (27), contradicts the minimality of C. Second, assume that C traverses only one hub, namely H. Then each of D_1, D_2 has fewer visits of H (and at most the same number of traversals) and we again obtain a contradiction to the minimality of C. This proves (25).

Since C is short, C cannot be contained in the frame F and therefore traverses a hub H. Then, by (25), the component $C \cap \overline{H}$ is a path, which we denote by Q_H. Its endvertices are two gates g, g' of H. If we replace Q_H in C by any g–g'-path in $\overline{H}$, we obtain a cycle, because otherwise $C \cap \overline{H}$ would have more than one component.

Let P_H be the (unique) g–g'-path in S_H, and assume first that $\ell(P_H) < 5\ell$. We replace in C the path Q_H by P_H and obtain a cycle C' such that $\ell(C') \leq \ell(C) + \ell(P_H) \leq 6\ell$. Thus together with (7), C' is a short cycle. Moreover, C' does not traverse H anymore as $C' \cap \overline{H} \subseteq S_H$. Thus, C' contradicts the minimal choice of C.

Second, assume that $\ell(P_H) \geq 5\ell$. By Claim 9, there is a g–g'-path P'_H in $\overline{H}$ with $\ell \leq \ell(P'_H) \leq 3\ell$. Thus, if we replace Q_H by P'_H in C, we obtain a cycle C' such that $\ell \leq \ell(P) \leq \ell(C') \leq \ell(C) + \ell(P'_H) \leq 4\ell$, which is the final contradiction to (7). □

4.4 The hitting set

We distinguish two cases: that F is a cycle ($U = \emptyset$) and $U \neq \emptyset$. Even if the first case could be transferred into the latter case, we found it useful to give a proof on its own.

Claim 12. *Unless there is a hitting set of at most $k-1$ edges, the frame F is not a cycle.*

Proof. Assume F to be a cycle. As shadows of hubs are trees, by Claim 7, every shadow of a hub is a path. In particular, the cycle F cannot lie in a single shadow. Thus, there are two distinct vertices u_1, u_2 in F that do not lie in the interior of any shadow (that is, if u_i is in a shadow, then it is an endvertex of the shadow).

Denote by P_1 and P_2 the two edge-disjoint u_1–u_2-paths in F. For $i = 1, 2$, we let $\overline{P}_i$ be the union of P_i and all hubs H so that $S_H \subseteq P_i$. Then

$$G = \overline{P}_1 \cup \overline{P}_2.$$

Indeed, any edge e of F is contained in $P_1 \cup P_2$. If $e \in E(G) \setminus E(F)$, then e lies in a hub, and every hub is contained in either $\overline{P}_1$ or in $\overline{P}_2$ as its shadow lies in either P_1 or in P_2.

Since hub closures are blocks in $\overline{P}_i$—the endvertices of their shadow-paths are cutvertices in $\overline{P}_i$—it follows from Claim 8 that every long cycle contains an edge of $\overline{P}_1$ and an edge of $\overline{P}_2$. More precisely, every long cycle can be decomposed into two u_1–u_2-paths—one in each $\overline{P}_i$.

Suppose that for $i = 1$ or for $i = 2$, there is a set X of at most $k-1$ edges that separates u_1 from u_2 in $\overline{P}_i$. Then, X meets every long cycle, since every such cycle contains a u_1–u_2-path in both $\overline{P}_1$ and $\overline{P}_2$. This means that X is a hitting set of size at most $k-1$, and we are done.

Thus, for $i = 1, 2$ there are k edge-disjoint u_1–u_2-paths $Q_1^i, \ldots, Q_k^i$ contained in $\overline{P}_i$. We combine them to k edge-disjoint cycles $Q_1^1 \cup Q_1^2, \ldots, Q_k^1 \cup Q_k^2$, each of which is long, by Claim 11, a contradiction to our assumption (6) that G does not contain k edge-disjoint long cycles. $\square$

We may assume from now on that F is not a cycle, and that therefore $U \neq \emptyset$. Since F is connected and has minimum degree at least 2, this implies that

$$F \text{ is the edge-disjoint union of } U\text{-paths.} \tag{28}$$

We distinguish two kinds of hubs: A hub H is a *vertex-hub* if $S_H \cap U \neq \emptyset$ and a *path-hub* otherwise. Oberserve that the shadow of a path-hub is completely contained in some U-path of F. Let $\mathcal{H}$ be the set of all vertex-hubs. A vertex-hub is shown in Figure 3, while the hub in Figure 6 is a path-hub.

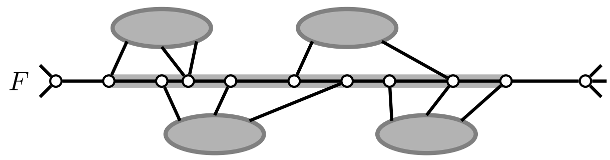

Figure 6: A path-hub consisting of four bridges, and its shadow (in grey).

For a hub H, let A_H be the set of gates of H and let $A^V = \bigcup_{H \in \mathcal{H}} A_H$.

Next, we give a bound from above for $\sum_{H \in \mathcal{H}} |A_H|$ for later use. We note that for every $g \in A^V$, the number of hub closures containing g is at most $d_F(g)$. Observe that every U-path of F contains at most two vertices of A^V. In addition, if P contains two vertices $g, g' \in A^V$ in its interior, then g, g' belong each to one vertex-hub only. This implies

$$\begin{aligned} \sum_{H \in \mathcal{H}} |A_H| &\leq \sum_{g \in A^V} d_F(g) \\ &= \sum_{g \in A^V \cap U} d_F(g) + \sum_{P:\ P \subseteq F \text{ is a } U\text{-path}} 2 \\ &\leq 2\text{ds}(F). \end{aligned}$$

Recalling (10) we obtain

$$\sum_{H \in \mathcal{H}} |A_H| \leq 168k \log k. \tag{29}$$

Consider a U-path P of F. If the shadow of a vertex-hub intersects P, then the intersection is either a path containing at least one endvertex of P, or the disjoint union of two paths each of which contains an endvertex of P. Thus at most one component of $P - \bigcup_{H \in \mathcal{H}} E(\overline{H})$ is a path of length at least 1. If there is such a component P', then let u_P, v_P be the endvertices of P'. Then $P' =$

$u_P P v_P$. Let $\mathcal{P}$ denote the set of all U-paths P of F such that $P - \bigcup_{H\in\mathcal{H}} E(\overline{H})$ is not edgeless. We note that

$$\textit{if } P \in \mathcal{P}, \textit{ then } u_P, v_P \in A^V \cup U.$$

For $P \in \mathcal{P}$, we define $\overline{P}$ to be the union of P' and all (path-)hubs H so that $S_H \subseteq P'$.

Next, we show

$$\textit{for any two distinct } A, B \in \mathcal{H} \cup \mathcal{P} \textit{ the graphs } \overline{A} \textit{ and } \overline{B} \textit{ are edge-disjoint and } \overline{A} \cap \overline{B} \subseteq A^V \cup \{u_P, v_P : P \in \mathcal{P}\}. \tag{30}$$

Indeed, this follows directly if both $A, B \in \mathcal{H}$, and also if both $A, B \in \mathcal{P}$, since U-paths in F meet only in U. If $A \in \mathcal{H}$ and $B \in \mathcal{P}$, then $u_B B v_B$ meets $\bigcup_{H\in\mathcal{H}} \overline{H}$ at most in $\{u_B, v_B\}$, by definition.

We claim that

$$G = \bigcup_{H\in\mathcal{H}} \overline{H} \cup \bigcup_{P\in\mathcal{P}} \overline{P} \tag{31}$$

To prove the claim, consider an edge $e \notin \bigcup_{H\in\mathcal{H}} E(\overline{H})$ of G. Assume first that e is contained in the closure of a path-hub L. The shadow of L then is contained in a U-path P of F, by (28). Since the shadow of L is edge-disjoint from $\bigcup_{H\in\mathcal{H}} \overline{H}$ this implies that $P \in \mathcal{P}$. Then $e \in E(\overline{L}) \subseteq E(\overline{P})$. Second, we have to consider the case when e is an edge of F that lies outside every hub shadow. Let P be the U-path of F containing e. Again we see that $P \in \mathcal{P}$ and trivially e is contained in $\overline{P}$. This proves (31).

Next we show

$$\textit{for every } P \in \mathcal{P}, \textit{ every cycle contained in } \overline{P} \textit{ is short.} \tag{32}$$

The graph $\overline{P}$ is the edge-disjoint union of path-hub closures and edges in F that lie outside every hub shadow. In particular, the path-hub closures contained in $\overline{P}$ are blocks in $\overline{P}$. Thus, any cycle contained in $\overline{P}$ lies completely in some path-hub closure, which only contains short cycles, by Claim 8.

We call $P \in \mathcal{P}$ *thick* if there are at least k edge-disjoint u_P–v_P-paths in $\overline{P}$, and *thin* otherwise. If P is thin, then there is a set $X_P \subseteq E(\overline{P})$ of at most $k-1$ edges separating u_P and v_P in $\overline{P}$, by Menger's theorem. As part of the hitting set we define X_p as the union of all X_P where $P \in \mathcal{P}$ is thin. By (10), we obtain

$$|X_p| = \sum_{P\in\mathcal{P}} |X_P| \le k \cdot \tfrac{1}{2}\mathrm{ds}(F) \le k \cdot 42k \log k.$$

We note that (32) implies that

$$\textit{in } G - X_p \textit{ every long cycle is edge-disjoint from } \overline{P} \textit{ for every thin } P \in \mathcal{P}. \tag{33}$$

Consider $H \in \mathcal{H}$. Applying Lemma 13 with $\overline{H}$ and A_H playing the roles of G and A, we obtain a set X_H of size at most $4|A_H|k$ that k-perfectly separates A_H in $\overline{H}$. Let $X_v = \bigcup_{H\in\mathcal{H}} X_H$. With (29) we find that

$$|X_v| \le 4k \sum_{H\in\mathcal{H}} |A_H| \le 4k \cdot 168k \log k = 672k^2 \log k.$$

We will show that $X = X_p \cup X_v$ is a hitting set for long cycles in G. We note first that

$$|X| = |X_p| + |X_v| \leq 42k^2 \log k + 672k^2 \log k = 714k^2 \log k \leq f(k, \ell).$$

Thus, if X is indeed a hitting set then the induction hypothesis (5) is proved.

Let $\mathcal{J}$ be the set of all graphs J such that either $J = \overline{P}$ for a thick $P \in \mathcal{P}$, or such that J is a component of $\overline{H} - X$ for some $H \in \mathcal{H}$.

Claim 13.

(i) *$g \sim_k g'$ in J for all $g, g' \in V(J) \cap (A^V \cup \{u_P, v_P : P \in \mathcal{P}\})$.*

(ii) *Distinct $J, J' \in \mathcal{J}$ are edge-disjoint, and their intersection $J \cap J'$ lies in $A^V \cup \{u_P, v_P : P \in \mathcal{P}\}$.*

(iii) *Every long cycle in $G - X$ is entirely contained in $\bigcup_{J \in \mathcal{J}} J$ and no long cycle is contained in a single $J \in \mathcal{J}$.*

Proof. Statement (i) holds as A_H is k-perfectly separated for every $H \in \mathcal{H}$ and if $J = \overline{P}$ for some thick $P \in \mathcal{P}$ then $u_P \sim_k v_P$ as P is thick and $\overline{P}$ disjoint from X.

Observe that (ii) follows from (30) as all $J \in \mathcal{J}$ are subgraphs of graphs in $\mathcal{H} \cup \mathcal{P}$ and two $J, J' \in \mathcal{J}$ that belong to the same vertex-hub H are disjoint by definition as components of $\overline{H} - X$.

To see (iii), consider a long cycle C. Since G is, by (31), the union of vertex-hub closures and all $\overline{P}$ for $P \in \mathcal{P}$, it follows that $G - X$ is contained in the union of all $J \in \mathcal{J}$ and all $\overline{P}$ for thin $P \in \mathcal{P}$. By (33), the cycle C is edge-disjoint from every $\overline{P}$, when $P \in \mathcal{P}$ is thin, which means that C is contained in the union of all $J \in \mathcal{J}$. Finally, C cannot be contained in any single $J \in \mathcal{J}$ as this is either a subgraph of a hub closure (recall Claim 8) or equal to $\overline{P}$ for some $P \in \mathcal{P}$ (recall (32)). $\square$

Suppose that $G - X$ contains a long cycle. Any long cycle C in $G - X$ decomposes by Claim 13 (iii) into paths $g_0 P_1 g_1, \ldots, g_s P_s g_1$ such that each P_i is contained in some $J_i \in \mathcal{J}$. Choose a long cycle C (and paths) such that the number s of paths P_i is minimal. That choice immediately guarantees that $J_i \neq J_{i+1}$ for all i (taken mod s).

Suppose that $J_i \cap J_j \neq \emptyset$ for $|i - j| > 1$. Since J_i is connected there is a C-path Q between two components of $C \cap J_i$. Then there are two cycles D_1, D_2 in $C \cup Q$ that contain Q. Let H be the hub such that $J_i \subseteq \overline{H}$. By Claim 8, C does not lie completely in $\overline{H}$. Thus, at least one of D_1 and D_2, D_1 say, also contains an edge outside $\overline{H}$. Thus, it follows from Claim 11 that D_1 is long. Then, however, D_1 contradicts the choice of C as it contains less paths P_i than C. Therefore, the J_i are all distinct.

Next, observe that, by Claim 13 (ii), every g_i either lies in A^V or in $\{u_P, v_P\}$ for some thick $P \in \mathcal{P}$. By Claim 13 (i), there are k edge-disjoint g_i–g_{i+1}-paths $P^i_1, \ldots, P^i_k$ in J_i for every $i = 0, \ldots, s$. Observe that the concatenation C_j of $P^1_j, \ldots, P^s_j$ is a cycle, which is long by Claim 11. Thus $C_1, \ldots, C_k$ are k edge-disjoint long cycles, which is the final contradiction to (6). Thus, the set X is indeed a hitting set for the long cycles in G. $\square$

References

[1] *ISGCI*, `http://www.graphclasses.org/smallgraphs.html`.

[2] E. Birmelé, J.A. Bondy, and B. Reed, *The Erdős-Pósa property for long circuits*, Combinatorica **27** (2007), 135–145.

[3] H. Bruhn, F. Joos, and O. Schaudt, *Long cycles through prescribed vertices have the Erdős-Pósa property*, arXiv:1412.2894 (2014).

[4] I.J. Dejter and V. Neumann-Lara, *Unboundedness for generalized odd cyclic transversality*, Colloq. Math. Soc. János Bolyai **52** (1987), 195–203.

[5] R. Diestel, *Graph theory*, fourth ed., Springer, Heidelberg, 2010.

[6] P. Erdős and L. Pósa, *On the maximal number of disjoint circuits of a graph*, Publ. Math. Debrecen **9** (1962), 3–12.

[7] S. Fiorini and A. Herinckx, *A tighter Erdős-Pósa function for long cycles*, J. Graph Theory **77** (2014), 111–116.

[8] T. Huynh, F. Joos, and P. Wollan, *A unified Erdős-Pósa theorem for constrained cycles*, arXiv:1605.07082 (2016).

[9] N. Kakimura, K. Kawarabayashi, and D. Marx, *Packing cycles through prescribed vertices*, J. Combin. Theory (Series B) **101** (2011), 378–381.

[10] K. Kawarabayashi and Y. Kobayashi, *Edge-disjoint odd cycles in 4-edge-connected graphs*, J. Combin. Theory (Series B) **119** (2016), 12–27.

[11] C.-H. Liu, *Packing and Covering immersions in 4-Edge-Connected Graphs*, arXiv:1505.00867 (2015).

[12] W. Mader, *Existenz n-fach zusammenhängender Teilgraphen in Graphen genügend grosser Kantendichte*, Abh. Math. Sem. Univ. Hamburg **37** (1972), 86–97.

[13] ———, *Über die Maximalzahl kantendisjunkter A-Wege*, Arch. Math. (Basel) **30** (1978), 325–336.

[14] F. Mousset, A. Noever, N. Škorić, and F. Weissenberger, *A tight Erdős-Pósa function for long cycles*, arXiv:1603.07588 (2016).

[15] M. Pontecorvi and P. Wollan, *Disjoint cycles intersecting a set of vertices*, J. Combin. Theory (Series B) **102** (2012), 1134–1141.

[16] J.-F. Raymond, I. Sau, and D. M. Thilikos, *An edge variant of the Erdős–Pósa property*, Discrete Math. **339** (2016), 2027–2035.

[17] J.-F. Raymond and D. Thilikos, *Recent techniques and results on the Erdős-Pósa property*, arXiv:1603.04615 (2016).

[18] N. Robertson and P. Seymour, *Graph minors. V. Excluding a planar graph*, J. Combin. Theory (Series B) **41** (1986), 92–114.

[19] M. Simonovits, *A new proof and generalizations of a theorem of Erdös and Pósa on graphs without $k+1$ independent circuits*, Acta Math. Acad. Sci. Hungar. **18** (1967), 191–206.

[20] C. Thomassen, *On the presence of disjoint subgraphs of a specified type*, J. Graph Theory **12** (1988), 101–111.

Version July 7, 2016

Henning Bruhn `<henning.bruhn@uni-ulm.de>`
Matthias Heinlein `<matthias.heinlein@uni-ulm.de>`
Institut für Optimierung und Operations Research
Universität Ulm, Ulm
Germany

Felix Joos `<f.joos@bham.ac.uk>`
School of Mathematics
University of Birmingham, Birmingham
United Kingdom