Frames, A-paths and the Erdős-Pósa property

Henning Bruhn[*] Matthias Heinlein Felix Joos[†]

Abstract

A key feature of Simonovits' proof of the classic Erdős-Pósa theorem is a simple subgraph of the host graph, a *frame*, that determines the outcome of the theorem. We transfer this frame technique to A-paths. With it we deduce a simple proof of Gallai's theorem, although with a worse bound, and we verify the Erdős-Pósa property for long and for even A-paths. We also show that even A-paths do not have the edge-Erdős-Pósa property.

1 Introduction

If a graph is far away from being a tree, is there a simple certificate for that — for instance a large number of vertex-disjoint cycles? The classic theorem of Erdős and Pósa asserts that, yes, there is such a certificate:

Theorem 1 (Erdős and Pósa [9]). *For every positive integer k, every graph either contains k vertex-disjoint cycles or a set of $O(k \log k)$ vertices that meets every cycle.*

The proof of Erdős and Pósa, while not long, is somewhat indirect and in particular rests on a result about A-paths of Gallai. (We will come back to that.) Simonovits [26] gave a different and cleaner proof that works in two steps. First Simonovits showed that a large cubic multigraph always contains many disjoint cycles. (Whenever we write "disjoint" we mean "vertex-disjoint".)

Lemma 2 (Simonovits [26]). *Every cubic multigraph with at least $(4 + o(1))k \log k$ vertices contains k disjoint cycles.*

In the second step, Simonovits considered a maximal subgraph F of the graph G with all degrees between 2 and 3. Then, if F has many vertices of degree 3, the lemma yields k disjoint cycles, and if not then the vertices of degree 3 will (almost) be a *hitting set*; that is, a set meeting all cycles.[1]

We think of the graph F in Simonovits' argument as a *frame*, a subgraph that essentially captures all target objects, the cycles, but has a much simpler structure. In particular, the frame alone determines whether we find k disjoint cycles or a hitting set, both of which can be (essentially) obtained directly from the frame.

[*]Partially supported by DFG, grant no. BR 5449/1-1

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

[1]We gloss over some technicalities here that, however, are not hard to resolve. In particular, components of F that consists of single cycles need to be taken care of, as well as cycles that meet F in a single vertex.

Pontecorvi and Wollan [21] extended the frame argument to capture also *A-cycles*, cycles that each contain at least one vertex of a fixed set A. Fiorini and Herinckx [10] used the technique in a different way to treat *long cycles*. Bruhn, Joos and Schaudt [3] further modified the technique so that it also applies to long A-cycles. Finally, Mousset, Noever, Škorić, and Weissenberger [20] again refined the frame argument in order to achieve an asymptotically tight bound (up to a constant factor) on the size of the hitting set for long cycles. A frame also plays an important role in our article [10] on edge-disjoint long cycles; see Section 6.

We demonstrate here that a simple frame argument also works for Erdős-Pósa type theorems about different variants of A-paths. In this case, trees with their leaves in A will play the role of a frame. We see the main merit of this article in discussing the frame argument and in showing how it can be used in a different context. Along the way, we will obtain several new results.

Gallai discovered that *A-paths*, paths with first and last vertex but no interior vertex in A and at least one edge, behave in a quite similar way as the cycles in Erdős and Pósa's theorem.

Theorem 3 (Gallai [11])**.** *For every positive integer k, every graph G and every set $A \subseteq V(G)$, the graph G either contains k disjoint A-paths or a vertex set X of size $|X| \leq 2k-2$ that meets every A-path.*

A class of objects, such as A-paths or cycles, has the *Erdős-Pósa property* if they satisfy a theorem similar to Theorems 1 and 3: namely that in every graph and for every k, there are either k disjoint such objects or that there is a set of vertices that meets each target object and whose size is bounded by function that depends only on k.[2]

Thus, cycles and A-paths have the Erdős-Pósa property, but also many other graph classes such as even cycles [27], graphs that contract to a fixed planar graph [23], or different variants of A-paths [6, 12, 15, 29].

Very general types of A-paths can be realised by labeling the edges of the graph with elements from an (abelian) group Γ. There are at least two ways to do that. In the simpler way, *undirected group labelings*, every edge e receives a label $\gamma(e) \in \Gamma$. A path P is then said to be *non-zero* if the sum of its edge labels is non-zero in Γ. If, for instance, the group Γ is $\mathbb{Z}_2$ and every edge receives a label of 1, then the non-zero paths are simply the paths of odd length. With respect to undirected labelings, Wollan [29] proved:

Theorem 4 (Wollan [29])**.** *For every graph G, vertex set $A \subseteq V(G)$, abelian group Γ and undirected Γ-edge labeling of G, there are either k disjoint non-zero A-paths or a vertex set of size $O(k^4)$ that meets every non-zero A-path.*

Chudnovsky et al. [6] investigated *directed group labelings*. In contrast to undirected labelings, an edge e of a path is counted with weight $\gamma(e)$ if it is traversed in direction of a fixed reference orientation, and with weight $-\gamma(e)$ otherwise. Chudnovsky et al. prove a result similar to Theorem 4 but with a

[2]That the notion is somewhat vague is on purpose. We want it to cover types of subgraphs, such as cycles, even cycles and so on, as well as subgraphs with additional structure, such as A-paths. To squeeze all these into one formally correct notion seems too much effort for the little benefit.

much smaller hitting set, namely one of size at most $2k-2$, which is clearly optimal.

Both results have interesting consequences: odd A-paths have the Erdős-Pósa property (Geelen et al. [12]), and so do *A–B–A-paths*, i.e. A-paths that each meet some vertex from a second vertex set B (see also Kakimura, Kawarabayashi, and Marx [15]).

Other extensions of Gallai's theorem concern directed A-paths (Kriesell [17]) and edge-disjoint A-paths (Mader [18]).

In this article, we apply the frame argument to show that A-paths (Section 2), long A-paths (Section 3), even A-paths (Section 4) as well as certain trees with their leaves in A (Section 5) have the Erdős-Pósa property. We also discuss types of A-paths that do not have the Erdős-Pósa property. These are A-paths with certain, more complicated, modularity constraints (Section 4) as well as directed even A-paths or directed A–B–A-paths (Section 7). In Section 6 we turn to the edge-version of the Erdős-Pósa property. In particular, we give a simple proof of Mader's theorem about edge-disjoint A-paths (but with a worse bound), and we show that neither even A-paths nor A–B–A-paths have the edge-Erdős-Pósa property. To the best of our knowledge this is the first example of a class that possesses the vertex-Erdős-Pósa property but not the edge-Erdős-Pósa property.

In the appendix, we provide a summary of several Erdős-Pósa property results.

2 Gallai's theorem

While Gallai's original proof of his theorem was quite technical, there are very nice proofs that reduce the problem to matchings and in particular to the Tutte-Berge formula; see Schrijver [25].

We give an alternative proof of Gallai's theorem, albeit with a slightly worse bound on the size of the hitting set. Gallai's bound of $2k-2$, on the other hand, is optimal as can be seen by considering a complete graph on $2k-1$ vertices, all in A.

Still, because the proof is quite simple and because it serves as a model for the other results later, we find it worth the effort. First, we set up the frame, the structure that captures the essence of the A-paths in the graph. The frame could not be much simpler: it is a tree with all its leaves in A. The following easy lemma plays the same role as Lemma 2 in Simonovits' proof.

Lemma 5. *Every subcubic tree with p leaves contains $\lfloor \frac{p}{2} \rfloor$ disjoint leaf-to-leaf paths.*

Proof. Let T be a subcubic tree with p leaves. We proceed by induction on the number of leaves. By contracting the edges incident with vertices of degree 2, we obtain a new tree whose leaf-to-leaf paths are in direct correspondence to leaf-to-leaf paths in T. We thus may assume that T contains no vertices of degree 2.

If T has at most three leaves, the statement is obviously true. Assume T to have at least four leaves, and pick a root r. Choose a vertex t of degree 3 that is farthest from r. Then t is adjacent to two leaves ℓ_1 and ℓ_2 and one non-leaf s (as T has at least four leaves). We remove the path $P = \ell_1 t \ell_2$ from T and obtain a new tree T'. As s has degree 3 in T, this implies that the sets of leaves of T and of T' differ only in ℓ_1 and ℓ_2. Inductively, we obtain $\lfloor \frac{p-2}{2} \rfloor$ disjoint leaf-to-leaf paths in T' and thus in T. Together with P we find the desired paths. $\square$

We prove Gallai's theorem with a hitting set of size at most $4k$.

Proof of Theorem 3 with hitting set of size at most $4k$. Let $F \subseteq G$ be a forest maximal under inclusion such that

- F is subcubic with no isolated vertices; and
- every leaf of F lies in A, and every vertex in $A \cap V(F)$ is a leaf of F.

Let c be the number of components of F. If F contains $2k+c$ vertices of A, then by applying Lemma 5 to each component of F we obtain k disjoint leaf-to-leaf paths in F, each of which is an A-path.

Thus, assume that $|A \cap V(F)| < 2k + c$. Let X be the union of $A \cap V(F)$ together with all vertices of degree 3 in F. In every subcubic tree, the number of vertices of degree 3 equals the number of leaves minus 2. Thus

$$|X| \leq 2|A \cap V(F)| - 2c < 4k.$$

We claim that X meets every A-path in G. Suppose that P is an A-path that is disjoint from X. It has to intersect F since otherwise $F \cup P$ would contradict the choice of F. Let P' be the initial segment of P that is an A–F-path, and observe that P' has length at least 1 since every vertex in $A \cap V(F)$ lies in X, and that P' does not end in a leaf of F since every leaf of F lies in X as well. Moreover, P' does not end in a vertex of degree 3 either, since these also all lie in X. The path P' therefore ends in a vertex of degree 2, which means that $F \cup P$ satisfies the same conditions as F and thus contradicts the choice of F. $\square$

3 Long A-paths

For a fixed positive integer ℓ a cycle is *long* if its length is at least ℓ. Long cycles have the Erdős-Pósa property [1, 10, 20, 23]. Analogously, we say a path is *long* if its length is at least ℓ. Here we give a short proof that long A-paths also have the Erdős-Pósa property.

Theorem 6. *For every positive integer k, every graph G and every set $A \subseteq V(G)$, the graph G either contains k disjoint long A-paths or a vertex set X of size $|X| \leq 4k\ell$ that meets every long A-path.*

For the proof we adapt, in a fairly straightforward way, a technique of Fiorini and Herinckx [10] and combine it with the forest-frames of the previous section.

Proof. Let $F \subseteq G$ be a forest maximal under inclusion such that

- F is subcubic with no isolated vertices;
- every leaf of F lies in A, and every vertex in $A \cap V(F)$ is a leaf of F; and
- every A-path in F is long.

Let c be the number of components of F. If F contains $2k + c$ vertices of A, then Lemma 5 applied to each component of F yields k disjoint leaf-to-leaf paths in F, each of which is a long A-path.

Thus, assume that $|A \cap V(F)| < 2k + c$. Let U be the set of vertices of degree 3 in F. We extend U to a vertex set X as follows: for every vertex $v \in A \cap V(F)$, we add all vertices in F at distance at most $\ell - 1$ from v in F.

To bound $|X|$ from above, assign every vertex of degree 2 and 3 in F at distance at most $\ell - 1$ to a closest leaf (which is in A by definition) and in case there are multiple choices decide according to some ordering of A. Observe that this results in an assignment such that all vertices that are assigned to a specific leaf induce a subtree of F. Now, let T be such a subtree, and pick a root $r \in A \cap V(T)$. By induction on the number m of vertices of degree 3 in T we prove that $|V(T)| \leq \ell + (\ell - 1) \cdot m$. If $m = 0$, we have $|V(T)| \leq \ell$. Otherwise, let p be a leaf in T distinct from r. The p–r-path P contains a vertex of degree 3 in T. Starting from p, let u be the first vertex of degree 3 in P and let v be the predecessor of u in P. Deleting pPv yields a subtree where u has no longer degree 3. Using the induction hypothesis leads to the desired result. Hence

$$\begin{aligned} |X| &\leq \ell|A \cap V(F)| + \ell|U| \\ &= \ell(|A \cap V(F)| + |A \cap V(F)| - 2c) \\ &< 4k\ell \end{aligned}$$

where we have also used that the number of leaves minus 2 equals the number of vertices of degree 3 in every component of F.

Suppose that there is a long A-path P that is disjoint from X. By the maximal choice of F, the path P meets F. Let P' be the initial segment of P that is an A–F-path. Note that P' ends in a vertex that has degree 2 in F, as $(V(F) \cap A) \cup U \subseteq X$. Finally, let Q be an A-path contained in $F \cup P'$ that starts in the A-vertex of P' and ends in $a \in A \cap V(F)$. Since X contains all vertices of F at distance at most $\ell - 1$ in F to a, it follows that Q is a long A-path. Thus $F \cup P'$ contradicts the choice of F. □

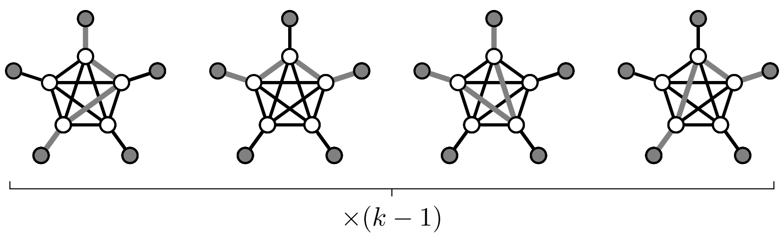

Figure 1: Example for long A-paths, with vertices in A in grey and $\ell = 4$; at most $k - 1$ disjoint long A-paths, yet every hitting set needs to contain $\ell - 1$ vertices from each component

How tight is the bound on the hitting set? The bound is likely not optimal, but has the right order of magnitude. To obtain a lower bound, take $k-1$ disjoint copies of a complete graph on $2\ell-3$ vertices whose vertices are matched to $2\ell-3$ vertices in A; see Figure 1 for an example with $\ell = 4$ and $k = 5$. Then in each component of the graph, there are no two disjoint long A-paths. A hitting set for long A-paths, on the other hand, contains at least $\ell - 1$ vertices from each component. Thus, there are no k disjoint long A-paths in the whole graph, while the smallest hitting set has size $(k-1)(\ell-1)$. We believe that in every graph there should be a hitting set of size at most $k\ell$. Our hitting set size of $4k\ell$ is only a bit off from that.

The construction is quite similar to one proposed by Fiorini and Herinckx [10] for long cycles.

4 Even A-paths

The results of Chudnovsky et al. or of Wollan (Theorem 4) discussed in the introduction imply in particular that odd A-paths have the Erdős-Pósa property. What about even A-paths? For cycles, there is a difference between even and odd cycles. The former have the Erdős-Pósa property, the latter do not. Interestingly, parity makes no difference for A-paths. Again, we apply the frame argument in the proof of the theorem.

Theorem 7. *For every positive integer k, every graph G and every set $A \subseteq V(G)$, the graph G either contains k disjoint even A-paths or a vertex set X of size $|X| \leq 10k$ that meets every even A-path.*

Proof. Let $F \subseteq G$ be a forest maximal under inclusion such that

- F is subcubic with no isolated vertices;
- every leaf of F lies in A, and every vertex in $A \cap V(F)$ is a leaf of F; and
- each component of F contains an even A-path.

Let c be the number of components of F. First assume that $|A \cap V(F)| \geq 4k + 2c$. If F has at least k components then, as each component contains an even A-path, there are k disjoint even A-paths. Thus, $c < k$.

Consider a component T of F. Then T is a tree, and its vertices split into two bipartition classes. The bipartition of T also partitions $A \cap V(T)$; let A_T be the one class of $A \cap V(T)$ that is not smaller than the other one. (If both are equal sized, pick one.)

Each vertex in $A \cap V(T)$ is a leaf of T. Delete the ones in $A \setminus A_T$ and iteratively their neighbours until the resulting tree T' has all its leaves in A_T (while keeping $A_T \subseteq V(T)$).

An application of Lemma 5 yields

$$\left\lfloor \frac{|A_T|}{2} \right\rfloor \geq \left\lfloor \frac{\lceil |A \cap V(T)|/2 \rceil}{2} \right\rfloor \geq \frac{|A \cap V(T)|}{4} - \frac{1}{2}$$

disjoint A_T-paths in T, and thus as many disjoint even A-paths.

Summing over all components we find at least

$$\sum_T \frac{|A \cap V(T)|}{4} - \frac{1}{2} = \frac{|A \cap V(F)|}{4} - \frac{c}{2} \geq k$$

disjoint even A-paths.

Second assume that $|A \cap V(F)| < 4k + 2c$. Let X be the union of $|A \cap V(F)|$ together with the set of vertices of degree 3 in F. Since the number of leaves minus 2 equals the number of vertices of degree 3 in a non-trivial subcubic tree it follows that

$$|X| \le 2|A \cap V(F)| - 2c \le 8k + 2c \le 10k,$$

where we used that $c < k$.

We claim that X is a hitting set for even A-paths. Suppose that X fails to meet some even A-path P. Then, P cannot be disjoint from F as otherwise $F \cup P$ would be better choice for F. Thus, P meets F. Let u be the first vertex of P in F (considered from some endvertex v of P). By definition of X, it follows that $u \notin A$ and that u does not have degree 3 in F. Therefore, u has degree 2 in F, and again $F \cup vPu$ contradicts the maximal choice of F. $\square$

By combining the proof techniques of Theorems 6 and 7, one may readily deduce that also *long even A-paths* have the Erdős-Pósa property.

What can we say about the size of the hitting set? While the bound in the theorem is not optimal, it turns out that the hitting set sometimes needs to be larger than $2k - 2$, the optimal bound for A-paths without parity constraints. To see this we can use the construction for long A-paths, where we set ℓ to 4; see Figure 1. The graph in that construction does not contain any A-paths of length 2; that is, every even A-path has length at least 4 and is thus long. Consequently, the graph does not contain k disjoint even A-paths, but at least $3k - 3$ vertices are necessary to meet every even A-path. We conjecture that a hitting set never needs more than $3k - 3$ vertices.

Cycles of quite general modularity constraints have the Erdős-Pósa property. This is the case, for instance, for cycles of length congruent to 0 modulo m, for every fixed positive integer m (Thomassen [27]); and it is also the case for cycles of a non-zero length modulo m, whenever m is odd (Wollan [30]). While non-zero A-paths are quite well covered by the results of Chudnovsky et al. and of Wollan (see Introduction), not much seems to be known about *zero* A-paths, A-paths of zero length modulo m, or more generally, with weight $\gamma(P) = 0$ for some directed or undirected group labeling γ of the edges with elements of an abelian group. Is there a counterpart of Theorem 4 for zero A-paths?

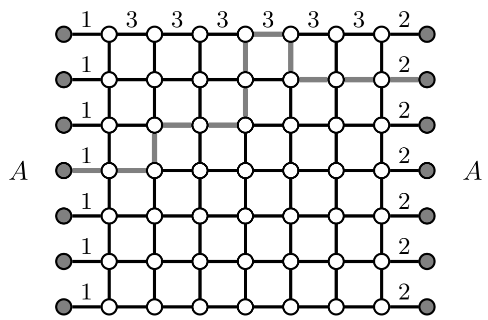

Figure 2: All unlabeled edges have length 6; an A-path of length divisible by 6 in grey.

No, it turns out: already A-paths of length divisible by 6 fail to have the Erdős-Pósa property. To see this, consider the graph in Figure 2. It consists of a subdivision of an $10r \times 10r$-grid, whose left side is matched to r vertices in A, and whose right side is linked by disjoint paths of length 2 to r different vertices in A. Finally, the top edges of the grid are subdivided to have length 3 each, while the other, unmarked edges of the grid are subdivided to have a length of 6 each.

Every A-path with both its endvertices on the left will have a length of $2+3s$ for some integer s, while any A-path with both endvertices on the right has a length of $4+3t$, for some integer t. Neither of these lengths is divisible by 6. Any A-path with a length divisible by 6 has to cross the grid from left to right, and in addition, needs to pick up an odd number of the subdivided top edges. Clearly, no two such paths can be disjoint. On the other hand, no hitting set can have size smaller than r.

We can generalise the construction to lengths divisible by m for other integers than $m=6$. Let $m>4$ be a composite number and let p be its smallest prime divisor. This implies $\frac{m}{p}>2$. Let $b=\frac{m}{p}$ and $c=m-\frac{m}{p}-1$. In the construction of Figure 2 we replace every length 3 by b, every length 2 by c and every length of an unmarked edge by m. Call the new graph $G(m,s)$ if the grid has size $s\times s$. Then, any A-path that has both endvertices on one side has length congruent to $\pm 2+k\cdot\frac{m}{p}\not\equiv 0 \pmod m$ because $\frac{m}{p}>2$. Any A-path that crosses the grid but avoids the top edges has a length of $m-\frac{m}{p}\not\equiv 0 \pmod m$. Thus, every A-path of a length divisible by m crosses from left to right and picks up at least one of the top edges. Again, there cannot be two such paths that are disjoint. Any hitting set, on the other hand, must contain a substantial part of the grid and thus has unbounded size.

We can even prove a more general statement on A-paths with modularity constraints:

Proposition 8. *For any composite integer $m>4$ and every $d\in\{0,\dots,m-1\}$, the A-paths of length congruent to d modulo m do not have the Erdős-Pósa property.*

Proof. We start with the graphs $G(m,s)$, the counterexamples for A-paths of length divisible by m, and then, depending on m and d, subdivide some of the edges incident with A. An A-path in any subdivision of $G(m,s)$ is *proper* if it starts on the left, intersects the top of the grid and ends on the right.

We will modify $G(m,s)$ in such a way that every proper A-path that intersects exactly one subdivided edge of the top of the grid has length congruent to d modulo m. That means, in particular, that every hitting set needs to have size proportional to s. On the other hand, we will show that no improper A-path can have length congruent to d modulo m. Since no two proper A-paths are disjoint, this will be enough to finish the proof.

If the equation $2x\equiv d \pmod m$ has a solution x, we modify the graph $G(m,s)$ as follows: replace every edge incident with a vertex in A by a path of length $x+1$. Then, every A-path in $G(m,s)$ of length ℓ corresponds to an A-path in the modified graph of length $\ell+2x\equiv\ell+d$, and vice versa. Therefore, with the same argument as above, we deduce that A-paths of length congruent to $d \pmod m$ in the modified graph are proper.

Next, assume that $2x\equiv d \pmod m$ does not have a solution, which implies that d is odd but m even. Thus, $p=2$. If $d\not\equiv\frac{m}{2}-2 \pmod m$, subdivide the edges in $G(m,s)$ incident with a vertex in A on the left of the grid such that they become paths of length $d+1$. Then, proper A-paths have length

$(d+1)+r\frac{m}{2}+(m-\frac{m}{2}-1)$, which is congruent to $d \pmod m$ for some odd r. Every improper A-path, in contrast, has a length congruent to $2(d+1)$, $2(d+1)+\frac{m}{2}$ (both ends on the left), $2(\frac{m}{2}-1)$, $2(\frac{m}{2}-1)+\frac{m}{2}$ (both ends on the right) or $(d+1)+(\frac{m}{2}-1)$ (no top intersection). Using that d is odd, m is even and $d \not\equiv \frac{m}{2}-2 \pmod m$, we see that none of these lengths are congruent to d.

Suppose now $d \equiv \frac{m}{2}-2$ (and d is odd and m is even). Now, subdivide the edges in $G(m,s)$ incident with a vertex in A on the right side of the grid such that they become paths of length $d+1$. Improper A-paths have length congruent to 2, $\frac{m}{2}+2$ (both ends on the left), $2(\frac{m}{2}-1+d)$, $2(\frac{m}{2}-1+d)+\frac{m}{2}$ (both ends on the right) or $1+\frac{m}{2}-1+d$. Using $d \equiv \frac{m}{2}-2$, $m>4$ and the parities of m and d, we see that none of these lengths are congruent to $d \pmod m$. $\square$

Intriguingly, the construction does not work if m is a prime or equal to 4. Does the Erdős-Pósa property hold in these cases? We do not know for the prime case but for $m=4$ this is indeed the case — the proof, however, needs quite a bit more work than the ones here; see [4].

In another simple generalisation, we add to the graph an undirected group labeling (see page 2). In many groups, in which there are suitable weights b,c to replace the lengths b,c, the construction can be adapted so that zero A-paths cannot have the Erdős-Pósa property.

In other groups this does not seem possible. The construction fails, for instance, when the underlying group is $\mathbb{Z}_2^\ell$ for some $\ell \in \mathbb{N}$. This is for a reason: the proof of Theorem 7 can be adapted so that it gives the Erdős-Pósa property for $\mathbb{Z}_2^\ell$-zero A-paths.

The situation seems to be more complicated for directed group labelings. In this setting, we currently can only construct a counterexample with the group $\mathbb{Z}^2$ but have been unable to do so for any finite group.

Arguably, the first Erdős-Pósa type result is Menger's theorem. Indeed, weakening it somewhat, we may reformulate Menger's theorem as: A–B-paths have the Erdős-Pósa property. Strikingly, and in contrast to A-paths, the property is lost once we impose parity constraints. For instance, neither even nor odd A–B-paths have the Erdős-Pósa property. This follows from an easy modification of the counterexample in Figure 2. We only have to replace the right part of A by B, and to adjust the lengths in the grid in such a way that every A–B-paths that avoids the top edges has the wrong parity and such that any A–B-path that traverses one of the top edges has the right parity.

5 Combs

The forest-frame technique is not only suited for different kinds of A-paths but may also be used for certain simple trees. One such example are *combs*.

Let us define an *elementary ℓ-comb*, for an integer $\ell \geq 1$, as the graph obtained from a path of length ℓ by adding a pendant edge to each internal vertex; see Figure 3. An *ℓ-comb* is any subdivision of an elementary ℓ-comb. Finally, for a given vertex set A, we say that an ℓ-comb is an *A-ℓ-comb* if all its leaves are in A and if every A-vertex in the comb is a leaf.

Theorem 9. *For any positive integer ℓ, there exists an integer c_ℓ such that following holds: For every positive integer k, every graph G, and every set*

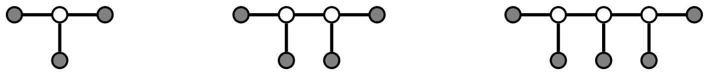

Figure 3: An A-2-comb, an A-3-comb and an A-4-comb

$A \subseteq V(G)$, the graph G either contains k disjoint A-ℓ-combs or a vertex set X of size $|X| \le c_\ell k$ that meets every A-ℓ-comb.

Proof sketch. As a frame we choose a $\subseteq$-maximal forest $F \subseteq G$ such that

- F is subcubic with no isolated vertices;
- every leaf of F lies in A, and every vertex in $A \cap V(F)$ is a leaf of F; and
- each component of F contains an A-ℓ-comb.

Similar as with the leaf-to-leaf paths in a tree, we may find $\lfloor \frac{4}{c_\ell} |A \cap V(T)| \rfloor$ disjoint A-combs in every component T of F, but at least one (for some positive integer c_ℓ). Thus with basically the same calculations as in the proofs of Theorems 6 and 7, we see that there are k disjoint A-ℓ-combs in F unless the set X, consisting of the vertices from A in F and of the vertices of degree 3 in F, has size smaller than $c_\ell k$. As before we may argue that X is a hitting set for A-ℓ-combs. Indeed, suppose there is an A-ℓ-comb C in $G - X$. Then, C contains an A–F-path that can be added to F, which results in a larger frame. □

6 Edge-versions and Mader's theorem

Theorem 1, the theorem of Erdős and Pósa, as well as Gallai's theorem (Theorem 3), both have an edge-version. The one of Gallai's theorem is due to Mader.

Theorem 10 (Mader [18])**.** *For every positive integer k, every graph G and every set $A \subseteq V(G)$, the graph G either contains k edge-disjoint A-paths or an edge set F of size $|F| \le 2k - 2$ that meets every A-path.*

Mader's proof is not short. Using a frame-like argument we give here a short proof but with a much worse bound. Again a tree will serve as frame:

Lemma 11 (Thomassen [28])**.** *Let T be a tree and let $A \subseteq V(T)$. Then T contains $\lfloor \frac{1}{2}|A| \rfloor$ edge-disjoint A-paths.*

The lemma is the finite version of a result of Thomassen [28, Theorem 8]. Thomassen writes that the finite version is an 'easy exercise'.

We now prove our weaker version of Mader's theorem.

Proposition 12. *For every positive integer k, every graph G and every set $A \subseteq V(G)$, the graph G either contains k edge-disjoint A-paths or an edge set F of size $|F| \le 2k \log_2 k$ that meets every A-path.*

Proof. We may assume that G is connected. Pick a spanning tree T of G. Now, if $|A| \ge 2k$ then, by Lemma 11, the graph G contains k edge-disjoint A-paths.

Thus, we may assume that $|A| < 2k$. Put $\mathcal{A}_0 = \{A\}$. Unless we find k edge-disjoint A-paths we construct for each $i = 1, \ldots, \lceil \log_2 |A| \rceil$ a partition $\mathcal{A}_i$

of A and an edge set X_i of size at most ik such that X_i meets every B–B'-path for distinct $B, B' \in \mathcal{A}_i$. Assume the construction to be achieved for $i-1$. Split each $B \in \mathcal{A}_{i-1}$ into two sets B_1,B_2 that differ in cardinality by at most 1, and let $\mathcal{A}_i$ be the set of all such B_1, B_2, and let B_1^* be the union of all such B_1, and let B_2^* be the union of all such B_2.

Apply Menger's theorem between B_1^* and B_2^*. If there are at most $k-1$ edges that separate B_1^* from B_2^* in $G - X_{i-1}$, then we take as X_i the union of these edges with X_{i-1}. Since X_{i-1} separates, by induction, B from B' for any two distinct $B, B' \in \mathcal{A}_{i-1}$, it follows that X_i separates any two sets in $\mathcal{A}_i$. If there is no edge set of size at most $k-1$ that separates B_1^* from B_2^*, then there are k edge-disjoint B_1^*–B_2^*-paths, each of which is an A-path.

Note that for $s = \lceil \log_2 |A| \rceil$ each set in $\mathcal{A}_s$ is a singleton. Thus, X_s meets every A-path. Its size is $|X_s| \leq \lceil \log_2 |A| \rceil \cdot k < \lceil \log_2(2k) \rceil k$. □

Arguably, the key argument differs markedly from the other arguments in this article, and should perhaps not be a called a frame argument. Indeed, as a characteristic of a frame we stated at the beginning that the frame determines the outcome: either it is large, with respect to some suitable measure, and then yields k disjoint target objects, or it is small and delivers a hitting set. In the proof of Proposition 12, the frame, the spanning tree, only gets us halfway: if it is large, that is, if it contains many vertices from A, then we find k edge-disjoint A-paths, but if it is small (not many A-vertices), then both outcomes are still possible.

Why is that so? Did we not use the 'right' frame? Perhaps there simply is no 'right' frame. Indeed, edge-versions in this context seem to be generally more complicated. To see this, let us come back to the theorem of Erdős and Pósa.

Let us say that a class of graphs (or more generally objects) $\mathcal{F}$ has the *edge-Erdős-Pósa property* if for every integer k, there is a number $f(k)$ such that every graph either contains k *edge-disjoint* subgraphs each isomorphic to some element of $\mathcal{F}$ or an *edge set* Z of size at most $f(k)$ that meets every subgraph contained in $\mathcal{F}$.

Simonovits' proof of Theorem 1 can be modified so that it yields the edge-Erdős-Pósa property for cycles, and Mader's theorem shows that A-paths have it, too. Some more classes are known to have the edge-Erdős-Pósa property, but not as many as are known to have the ordinary, vertex-version, property. In some rare cases, the edge-property can be deduced from the vertex-property. This is, for instance, the case for even cycles. That even cycles have the ordinary Erdős-Pósa property seems to have been observed first by Neumann-Lara:

Theorem 13 (Neumann-Lara, see [7] or [27])**.** *Even cycles have the vertex-Erdős-Pósa propery. That is, there is a function f such that for every positive integer k every graph G either contains k disjoint even cycles or there is a vertex set X of size $|X| \leq f(k)$ such that $G - X$ does not contain any even cycle.*

Chekuri and Chuzhoy [5] demonstrate that the size of the hitting set can be bounded by a function $f(k) = O(k \,\mathrm{polylog} k)$.

Theorem 14. *Even cycles have the edge-Erdős-Pósa property.*

We use here a technique that is similar to the one in the proof of Theorem 7.

Proof. Let f be a function as in Theorem 13. We may assume that every vertex is contained in some even cycle — otherwise we could delete the vertex without changing anything.

Assume first that G contains a vertex x of degree at least $6k$. Let c be the number of components of $G - x$, and consider a component K of $G - x$. Then there is an even cycle that meets K. Since the vertex set of any such even cycle is contained in $V(K) \cup \{x\}$, we therefore find at least c edge-disjoint even cycles, one for each component of $G - x$. Thus, we may assume that $c \leq k - 1$.

Subdivide every edge between x and K exactly once, and denote the set of subdividing vertices by A. In particular, $|A| = |N(x) \cap V(K)|$. Pick a spanning tree T of K (in the subdivided graph). The bipartition of T induces a bipartition of $A \cap V(K)$. Let A' be one of the two induced bipartition classes of A such that $|A'| \geq \lceil \frac{1}{2}|A \cap V(K)| \rceil$. Applying Lemma 11, we obtain $\lfloor \frac{1}{2}|A'| \rfloor$ many edge-disjoint A'-paths contained in K, each of which is an even A-path. By replacing the first and last edge of such a path P by the edge between the second vertex of the path and x, and by the edge between the penultimate vertex and x, we obtain an even cycle. In this way we obtain pairwise edge-disjoint even cycles contained in $G[K + x]$. The number of these is at least

$$\left\lfloor \frac{\lceil \frac{1}{2}|N(x) \cap V(K)| \rceil}{2} \right\rfloor \geq \frac{1}{4}|N(x) \cap V(K)| - \frac{1}{2}.$$

Summing over all components K of $G - x$ we obtain at least

$$\sum_K \left(\tfrac{1}{4}|N(x) \cap V(K)| - \tfrac{1}{2} \right) = \tfrac{1}{4}|N(x)| - \tfrac{1}{2}c \geq k$$

edge-disjoint even cycles, where we have used that $|N(x)| \geq 6k$ and $c \leq k$.

It remains to consider the case when no vertex in G has degree at least $6k$. Since even cycles have the vertex-Erdős-Pósa property (Theorem 13), there is a vertex set X of size at most $f(k)$ such that $G - X$ does not contain any even cycle. Let F be the set of all edges incident with any vertex in X. Then F is an edge hitting set for even cycles of size $|F| \leq 6kf(k)$. □

Normally, the edge-property does not follow as easily. Long cycles, for instance, do have the edge-property but in contrast to the vertex-version, the proof is much longer and quite a bit more complicated [2]. While a frame argument is used, as in our proof of Mader's theorem, the frame is much weaker. If the frame is large, then k edge-disjoint long cycles are found, but if it is small, then more work is necessary and both outcomes are possible.

That we know less about the edge-Erdős-Pósa property becomes immediately apparent when we consider A-paths. It is an open problem, whether long A-paths have it.

Problem 15. *Do long A-paths have the edge-Erdős-Pósa property?*

We point out that another open problem would give an affirmative answer for long A-paths.

Problem 16. *Do long A-cycles have the edge-Erdős-Pósa property?*

Indeed, consider a graph G with a vertex set A, and a fixed length ℓ. Now, add a vertex s and link s to each $a \in A$ by $d(a)$ parallel paths each of length 2. In the resulting graph G' apply the edge-Erdős-Pósa property for long $\{s\}$-cycles, where we use a minimal length of $\ell + 4$. Then every long $\{s\}$-cycle contains a long A-path, and vice versa, every long A-path can be extended to a long $\{s\}$-cycle. Hitting sets may be translated in a similar fashion.

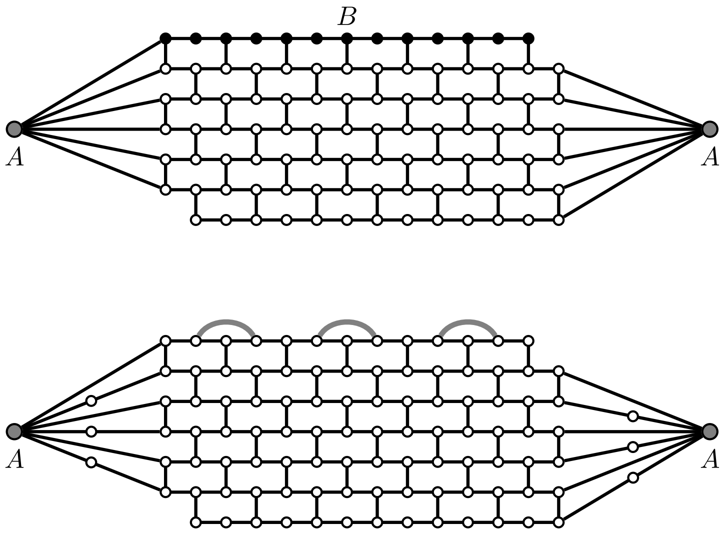

Figure 4: Counterexamples: neither A–B–A-paths nor even/odd A-paths have the edge-Erdős-Pósa property

If some class $\mathcal{F}$ does not have the vertex-Erdős-Pósa property, such as odd cycles, then it usually does not have the edge-Erdős-Pósa property either. This is because the counterexamples for the vertex-property are normally based on a grid plus some extra structure; these then often (always?) turn into counterexamples for the edge-Erdős-Pósa property if a wall is used instead. (A wall is the subcubic analogue of a grid; for a formal definition see for instance Robertson and Seymour [24]).

In contrast, classes $\mathcal{F}$ may have the vertex-Erdős-Pósa property but not the edge-Erdős-Pósa property. We show that here for A–B–A-paths, for even A-paths and for odd ones. To the best of our knowledge, such graph classes were not known before.

Consider the top graph in Figure 4: it consists of a wall with $10r \times 10r$ bricks together with two vertices in A, one of which is linked to the left side, while the other is linked to the right side of the wall. The top row of the wall makes up the vertex set B. Clearly, no two A–B–A-paths can be edge-disjoint.

At the same time, no hitting set of edges contains fewer than r edges. Indeed, let X be some edge set of fewer than r edges, let a_1 denote the left-hand side vertex in A and a_2 the right-hand side one. A wall of size $10r \times 10r$ contains $10r + 1$ disjoint vertical paths $P_0, \ldots, P_{10r}$, ordered according to their starting vertex in the top row. As $|X| < r$, there are at least two consecutive vertical paths P_i, P_{i+1} such that X is disjoint from $P_i \cup P_{i+1}$ and does not contain any of the two edges e_i, e'_i with endvertices in B between the two starting vertices of P_i and P_{i+1} either. Since there are, moreover, $10r$ edge-disjoint a_1–P_i-paths, and also that many edge-disjoint P_{i+1}–a_2-paths, the set X must miss at least one a_1–

P_i-path, Q_1 say, and at least one P_{i+1}–a_2-path, Q_2 say. Then $Q_1 \cup P_i \cup P_{i+1} \cup Q_2$ together with e_i, e'_i contains an A–B–A-path that avoids X. Thus, X cannot be a hitting set of edges. This shows that A–B–A-paths do not have the edge-Erdős-Pósa property.

The construction for odd (or even) A-paths is very similar, and shown in the bottom part of Figure 4. Here, by subdividing certain edges incident with the left or the right A-vertex, we make sure that every A-path that avoids the grey edges has even (resp. odd) length. If we define B as the set of endvertices of the grey edges, then every odd (resp. even) A-path is an A–B–A-path and we may argue as above.

A-combs, in the sense of Section 5, also fail to have the edge-Erdős-Pósa property. The counterexamples are very similar to the ones discussed in this section.

As in Section 4 we can generalise even or odd A-paths to zero A-paths with respect to some (directed or undirected) labeling of the edges with an abelian group Γ. While in the vertex-version the group might make a difference, this is not the case for the edge-Erdős-Pósa property. Indeed, the construction shown in the bottom part of Figure 4 turns into one for zero A-paths: we just label all edges incident with the left vertex in A with a non-zero group element $\mu \in \Gamma$, we label all grey edges with $-\mu$ and all other edges with 0 (in addition, if there is a reference orientation, then orient all edges from left to right and top to bottom).

7 Directed versions

Why do A–B–A-paths have the vertex-Erdős-Pósa property but not the edge-version? Because in the edge-version we can force the A–B–A-paths in examples such as in Figure 4 to cross the wall from left to right: as no path can return to its starting vertex, we can fix start and endvertex in the edge-version by only assigning two vertices to A. In the vertex-version this fails, as we could always put the two vertices in the hitting set. But if, instead, we replace the left and the right vertex in A by many vertices, then the A-paths can return to their starting side.

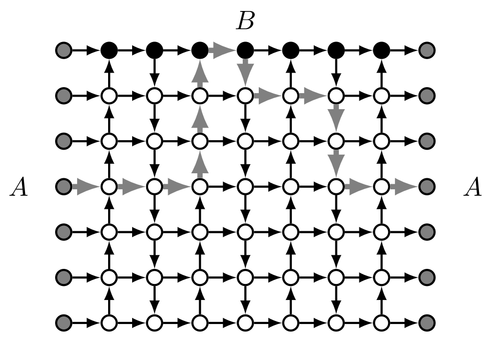

Figure 5: A directed A–B–A-path

It is intuitively clear that we can also enforce a direction of the A-paths in a digraph, and indeed, directed A–B–A-paths do not have the Erdős-Pósa property. To see this, consider the construction shown in Figure 5, where we again scale the size of the grid to be (much) larger than the size of any purported hitting set. Since every A–B–A-path needs to cross the grid from left to right, and needs to meet the top row as well, no two disjoint such paths are possible. Any (vertex) hitting set, however, will need to contain a number of vertices that grows with the size of the grid.

The construction easily transfers to the edge-Erdős-Pósa property if the grid is replaced by a wall as in Figure 4, and, in a similar way, also extends to even or odd directed A-paths.

References

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

[2] H. Bruhn, M. Heinlein, and F. Joos, *Long cycles have the edge-Erdős-Pósa property*, to appear in Combinatorica.

[3] H. Bruhn, F. Joos, and O. Schaudt, *Long cycles through prescribed vertices have the Erdős-Pósa property*, J. Graph Theory **87** (2018), 275–284.

[4] H. Bruhn and A. Ulmer, *Packing A-paths of length zero modulo four*, in preparation.

[5] C. Chekuri and J. Chuzhoy, *Large-treewidth graph decompositions and applications*, arXiv:1304.1577 (2013).

[6] M. Chudnovsky, J. Geelen, B. Gerards, L. Goddyn, M. Lohman, and P. Seymour, *Packing non-zero A-paths in group-labelled graphs*, Combinatorica **26** (2006), 521–532.

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

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

[9] P. Erdős and L. Pósa, *On independent circuits contained in a graph*, Can. J. Math. **7** (1965), 347–352.

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

[11] T. Gallai, *Maximum-minimum Sätze und verallgemeinerte Faktoren von Graphen*, Acta Math. Hungar. Acad. Sci. **12** (1961), 131–173.

[12] J. Geelen, B. Gerards, B. Reed, P. Seymour, and A. Vetta, *On the odd-minor variant of Hadwiger's conjecture*, J. Combin. Theory (Series B) **99** (2009), 20–29.

[13] M. Heinlein and A. Ulmer, *Directed cycles have the edge-Erdős-Pósa property*, arXiv:1802.05026 (2018).

[14] N. Kakimura and K. Kawarabayashi, *Packing directed circuits through prescribed vertices bounded fractionally*, SIAM J. Discrete Math. **26** (2012), 1121–1133.

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

[16] K. Kawarabayashi and S. Kreutzer, *The directed grid theorem*, arXiv:1411.5681 (2014).

[17] M. Kriesell, *Disjoint A-paths in digraphs*, J. Combin. Theory (Series B) **95** (2005), 168–172.

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

[19] L. Montejano and V. Neumann-Lara, *A variation of Menger's theorem for long paths*, J. Combin. Theory (Series B) **36** (1984), 213–217.

[20] F. Mousset, A. Noever, N. Škorić, and F. Weissenberger, *A tight Erdős-Pósa function for long cycles*, to appear in J. Combin. Theory (Series B).

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

[22] B. Reed, N. Robertson, P.D. Seymour, and R. Thomas, *Packing directed circuits*, Combinatorica **16** (1996), 535–554.

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

[24] ———, *Graph minors. XIII. The disjoint paths problem*, J. Combin. Theory (Series B) **63** (1995), 65–110.

[25] A. Schrijver, *A short proof of Mader's $\mathcal{S}$-paths theorem*, J. Combin. Theory (Series B) **82** (2001), 319–321.

[26] 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.

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

[28] ———, *Orientations of infinite graphs with prescribed edge-connectivity*, Combinatorica **36** (2016), 601–621.

[29] P. Wollan, *Packing non-zero A-paths in an undirected model of group labeled graphs*, J. Combin. Theory (Series B) **100** (2010), 141–150.

[30] ———, *Packing cycles with modularity constraints*, Combinatorica **31** (2011), 95–126.

A Overview of Erdős-Pósa results

In the table below, we summarise previous results, new results and open problems for different types of cycles and paths.

Class	**Vertex property**	**Edge property**
cycles	yes [9]	yes, e.g. [8, Ex. 9.6]
even cycles	yes, see [7] or [27]	yes, Theorem 14
odd cycles	no [7]	no*
cycles of length $\equiv 0 \pmod{m}$	yes [27]	open
cycles of length $\not\equiv 0 \pmod{m}$ with odd m	yes [30]	open
A-cycles	yes [21]	yes [21]
long cycles	yes [20]	yes [2]
long A-cycles	yes [3]	open
directed cycles	yes [22]	yes [13]
directed long cycles	yes [16]	open
directed A-cycles	no [14]	no*
A-paths	yes [11]	yes [18]
non-zero A-paths	yes [6, 29]	no, Figure 4
even A-paths	yes, Theorem 7	no, Figure 4
A-paths of length $\equiv 0 \pmod{m}$ with composite $m > 4$	no, Proposition 8	no, Proposition 8*
A-paths of length $\equiv 0 \pmod{4}$	yes [4]	no, Figure 4
A-paths of length $\equiv 0 \pmod{m}$ for m prime	open	no, Figure 4
long A-paths	yes, Theorem 6	open
long A–B-paths	yes [19]	open
even/odd A–B-paths	no, see Section 4	no, see Section 4*
A–B–A-paths	yes [15]	no, Figure 4
directed A-paths	yes [17]	open
directed long A–B-paths	yes [19]	open
directed A–B–A-paths	no, see Section 7	no*

*modify the counterexample by replacing the grid by a wall

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

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