Menger's theorem for a countable source set

Ron Aharoni and Reinhard Diestel

Dedicated to Paul Erdős
on the occasion of his 80th birthday

Paul Erdős has conjectured that Menger's theorem extends to infinite graphs in the following way: whenever A, B are two sets of vertices in an infinite graph, there exist a set of disjoint A–B paths and an A–B separator in this graph so that the separator consists of a choice of precisely one vertex from each of the paths. We prove this conjecture for graphs that contain a set of disjoint paths to B from all but countably many vertices of A. In particular, the conjecture is true when A is countable.

Introduction

If there is any conjecture in infinite graph theory whose fame has clearly transcended the boundaries of the field, it is the following infinite version of Menger's theorem, conjectured by Erdős:

Conjecture. (Erdős)
Whenever A, B are two sets of vertices in a graph G, there exist a set of disjoint A–B paths and an A–B separator in G so that the separator consists of a choice of precisely one vertex from each of the paths.

Here, G may be either directed or undirected and either finite or infinite, and 'disjoint' means 'vertex disjoint'. If G is finite, the statement is clearly a reformulation of Menger's theorem. A set of A–B paths together with an A–B separator as above will be called an *orthogonal paths/separator pair*.

We remark that the naïve infinite analogue to Menger's theorem which merely compares cardinalities, is considerably weaker and easy to prove. Indeed, consider any inclusion-maximal set $\mathcal{P}$ of disjoint A–B paths. If $\mathcal{P}$ can be chosen infinite then $\bigcup \mathcal{P}$, which is trivially an A-B separator, still has size only $|\mathcal{P}|$. If not, then choose $\mathcal{P}$ of maximal (finite) cardinality, and there is a simple reduction to the finite Menger theorem [5]. This was in fact first observed by Erdős, and seems to have inspired his above conjecture as the 'true' generalization of Menger's theorem.

Although Erdős's conjecture has been proved for countable graphs [2], a full proof still appears to be out of reach. However, no other conjecture in infinite graph theory has inspired as interesting a variety of partial or related results as this one; see [4] for a survey and list of references.

The main aim of this paper is to prove a lemma which, in addition to implying (with [2]) the results stated in the abstract, might play a role in an overall proof of the conjecture by induction on the size of G. Briefly, the lemma implies that if the conjecture is true for all graphs of size κ, where κ is any infinite cardinal, then it is true also for arbitrary graphs provided the source set A is no larger than κ. (In particular, we see that the conjecture holds for any graph if A is countable.) Now if $|A| = |G| = \lambda$ and the conjecture holds for all graphs of size $< \lambda$, the lemma enables us to apply the induction hypothesis to G with A replaced by its smaller subsets A'; we may then try to combine the orthogonal paths/separator pairs obtained between these A' and B to one between A and B. We must point out, however, that such a proof of Erdős's conjecture will by no means be straightforward, and it is not the only possible approach.

1. Definitions and statement of the main result

All the graphs we consider will be directed; undirected versions of our results can be recovered in the usual way by replacing each undirected edge with two directed edges pointing in opposite directions. An edge from a vertex x to a vertex y will be denoted by xy. When G is a graph, then $\overleftarrow{G}$ denotes the graph obtained from G by reversing all its edges.

Paths, likewise, will be directed, and we usually refer to them by their vertex sequence. If $P = x \dots y$ is a path and v, w are vertices on P in this order, then vPw denotes the subpath of P from v to w. Similarly, we write Pv and vP for initial and final segments of P, $P\mathring{v}$ for $Pv - v$, $\mathring{v}P$ for $vP - v$, and so on. If $Q = y \dots z$ is another path, and $P \cap Q = \{ y \}$, then PyQ denotes the path obtained by concatenating P and Q.

Let X, Y be sets of vertices in a graph. An X–Y *path* is a path from X to Y whose inner vertices are neither in X nor in Y. If x is a vertex, then a set of $\{ x \}$–Y paths that are disjoint except in x is an x–Y *fan*; the fan is *onto* if every vertex in Y is hit. Similarly, a set of X–y paths that are disjoint except in y is an X–y *fan*.

A *warp* is a set of disjoint paths. When $\mathcal{W}$ is a warp, we write $V[\mathcal{W}]$ for the set of vertices of the paths in $\mathcal{W}$, and $E[\mathcal{W}]$ for the set of their edges. Similarly, we write $in[\mathcal{W}]$ for the set of initial vertices of the paths in $\mathcal{W}$, and $ter[\mathcal{W}]$ for the set of their terminal vertices. For a vertex $x \in V[\mathcal{W}]$, we denote the path in $\mathcal{W}$ containing x by $Q_{\mathcal{W}}(x)$, or briefly $Q(x)$. For $x \notin V[\mathcal{W}]$, we put $Q_{\mathcal{W}}(x) := \{ x \}$. A warp consisting of A–B paths is an A–B *warp*. By $\overleftarrow{\mathcal{W}}$ we denote the warp in $\overleftarrow{G}$ consisting of the reversed paths from $\mathcal{W}$.*

* Clearly, $\overleftarrow{\overleftarrow{\mathcal{W}}} = \mathcal{W}$. We shall use this fact as an excuse to denote warps in $\overleftarrow{G}$, if they are introduced afresh rather than being obtained from a warp in G, by $\overleftarrow{\mathcal{W}}$ etc. straight away;

Let $G = (V, E)$ be a graph and $A, B \subseteq V$. Any such triple $\Gamma = (G, A, B)$ will be called a *web*. The web $(\overleftarrow{G}, B, A)$ is denoted by $\overleftarrow{\Gamma}$. An A–B warp $\mathcal{W}$ with $in[\mathcal{W}] = A$ is a *linkage* in Γ, and Γ is *linkable* if it contains a linkage.

A set $S \subseteq V$ *separates* A from B in G if every path in G from A to B meets S. Note that A and B may intersect, in which case clearly $A \cap B \subseteq S$.

A warp $\mathcal{W}$ in G is called a *wave* in Γ if $V[\mathcal{W}] \cap A = in[\mathcal{W}]$ and $ter[\mathcal{W}]$ separates A from B in G. The wave $\{(a) \mid a \in A\}$ is called the *trivial wave*. If $\mathcal{W}$ is a wave in Γ then $\Gamma/\mathcal{W}$ denotes the web

$$\Big(G - (A \backslash in[\mathcal{W}]) - (V[\mathcal{W}] \backslash ter[\mathcal{W}]) \ , \ ter[\mathcal{W}] \ , \ B \Big) .$$

In every web $\Gamma = (G, A, B)$ there is a wave $\mathcal{W}$ such that $\Gamma/\mathcal{W}$ has no non-trivial wave. (This is not difficult to see. If $\mathcal{W}_0$ is a wave in Γ and $\mathcal{W}_1$ is a wave in $\Gamma/\mathcal{W}_0$, then $\mathcal{W}_1$ defines a wave in Γ in a natural way: just extend its paths back to A along the paths of $\mathcal{W}_0$. This wave in Γ is 'bigger' than $\mathcal{W}_0$, and chains of waves in Γ with respect to this order tend to an obvious limit wave $\mathcal{W}$, which consists of the paths that are eventually in every wave of the chain. If the chain was maximal, then $\Gamma/\mathcal{W}$ has no non-trivial wave. See [2] for details.)

A wave $\mathcal{W}$ in Γ is a *hindrance* if $A \backslash in[\mathcal{W}] \neq \emptyset$; if Γ contains a hindrance, it is called *hindered*. Note that every hindrance is a non-trivial wave. The following was observed in [2]:

Erdős's conjecture is equivalent to the assertion that every unhindered web is linkable.

We are now in a position to state the main result proved in this paper. (For the reasons explained earlier, and because it is of a technical nature, we call it a lemma, not a theorem.)

Lemma 1. *Let $\Gamma = (G, A, B)$ be a web and $\mathcal{J}$ an A–B warp in G (possibly empty). If $|A \backslash in[\mathcal{J}]| > |B \backslash ter[\mathcal{J}]|$, then Γ is hindered.*

Lemma 1 will be proved in Sections 2 and 3. Our aim will be to turn the given warp $\mathcal{J}$, step by step, into a hindrance. This will require some alternating path techniques; the definitions and lemmas needed are given in Section 2. Section 3 is devoted to the main body of the proof of Lemma 1. In Section 4 we look at the implications of the lemma for Erdős's conjecture.

their reversals in G will then be denoted by $\mathcal{W}$. The idea here is to avoid the counter-intuitive practice of having a warp $\mathcal{W}$ in $\overleftarrow{G}$ and a resulting warp $\overleftarrow{\mathcal{W}}$ in G. This convention, if not its explanation, should help the reader to avoid any warps in his or her intuition when such things are discussed briefly in Section 4.

2. Alternating paths

Let $\Gamma = (G, A, B)$ be a web, and let $\mathcal{J}$ be an A–B warp in G. A finite sequence $P = x_0 e_0 x_1 e_1 \ldots e_{n-1} x_n$ of not necessarily distinct vertices x_i and distinct (directed) edges e_i of G will be called an *alternating path* (with respect to $\mathcal{J}$) if the following three conditions are satisfied:

(i) for every $i < n$, either $e_i = x_i x_{i+1} \in E(G) \backslash E[\mathcal{J}]$ or $e_i = x_{i+1} x_i \in E[\mathcal{J}]$;

(ii) if $x_i = x_j$ for $i \neq j$ then $x_i \in V[\mathcal{J}]$;

(iii) for every i, $0 \leqslant i < n$, if $x_i \in V[\mathcal{J}]$ then $\{ e_{i-1}, e_i \} \cap E[\mathcal{J}] \neq \emptyset$.

All the alternating paths we consider in this section will be alternating paths in G with respect to $\mathcal{J}$. Note that, by (iii) above, an alternating path starting at a vertex of $\mathcal{J}$ has its first edge in $\mathcal{J}$. As the edges of an alternating path are pairwise distinct, it can visit any given vertex at most twice, and this happens in essentially only two ways: if $x_i = x_j$ for $i < j < n$, then $x_i \in V[\mathcal{J}]$ by (ii), so by (iii)

either $e_{i-1}, e_j \in E[\mathcal{J}]$ and $e_i, e_{j-1} \notin E[\mathcal{J}]$ (Fig. 1 left)
or $e_i, e_{j-1} \in E[\mathcal{J}]$ and $e_{i-1}, e_j \notin E[\mathcal{J}]$ (Fig. 1 right).

FIGURE 1. Two alternating paths with respect to $\mathcal{J}$

Note that initial segments of alternating paths are again alternating paths, but final segments need not be. Finally, an ordinary path which avoids $\mathcal{J}$ or meets it only in its last vertex is trivially an alternating path.

There are analogous alternating versions of the notions of X–Y path, X–y fan and so on.

Lemma 2.1. *If $a \in A \backslash in[\mathcal{J}]$ and $b \in B \backslash ter[\mathcal{J}]$, and if $P = a \ldots b$ is an alternating path with respect to $\mathcal{J}$, then G contains an A–B warp $\mathcal{J}'$ such that $in[\mathcal{J}'] = in[\mathcal{J}] \cup \{ a \}$ and $ter[\mathcal{J}'] = ter[\mathcal{J}] \cup \{ b \}$ and $\{ Q \in \mathcal{J} \mid P \cap Q = \emptyset \} \subseteq \mathcal{J}'$.*

Proof. Consider the graph on $V[\mathcal{J}] \cup V(P)$ whose edge set is the symmetric difference of $E[\mathcal{J}]$ and $E(P)$. The (undirected) components of this graph are all finite. Considering their vertex degrees, we see that they are either A–B paths or cycles avoiding $A \cup B$ (possibly trivial). The assertion follows. $\square$

Lemma 2.2. *Let $P_1 = x_0e_0 \dots e_{n-1}x_n$ and $P_2 = y_0f_0 \dots f_{m-1}y_m$ be alternating paths. If $x_n = y_0$, then there exists an alternating path P_3 from x_0 to y_m such that $V(P_3) \subseteq V(P_1) \cup V(P_2)$ and $E(P_3) \subseteq E(P_1) \cup E(P_2)$.*

Proof. Let $i \leqslant n$ be minimal such that there exists a $j \leqslant m$ with the following two properties:

(i) $x_i = y_j$;

(ii) if $x_i \in V[\mathcal{J}]$, then either $e_{i-1} \in E[\mathcal{J}]$ or $f_j \in E[\mathcal{J}]$.

(Note that such an i exists, because $x_n = y_0$ and P_2 is an alternating path. Moreover, j is easily seen to be unique.) Then $x_0e_0 \dots e_{i-1}x_if_j \dots f_{m-1}y_m$ is an alternating path as desired. $\square$

3. Proof of the main lemma

We now prove Lemma 1. As in the lemma, let $\Gamma = (G, A, B)$ be a web, and let $\mathcal{J}$ be an A–B warp in Γ. Let us write

$$A_1 := \mathit{in}\,[\mathcal{J}] \qquad \text{and} \qquad A_2 := A \backslash A_1$$

and

$$B_1 := \mathit{ter}\,[\mathcal{J}] \qquad \text{and} \qquad B_2 := B \backslash B_1,$$

and put $\kappa := |B_2|$. We assume that $|A_2| > \kappa$, and construct a hindrance $\mathcal{W}$ in Γ. Again, all the alternating paths considered in this section will be alternating paths in G with respect to $\mathcal{J}$, unless otherwise stated.

Let us quickly dispose of the case when κ is finite. Assume that κ is minimal such that the lemma fails. By Lemma 2.1 and the minimality of κ, there is no alternating path from A_2 to B_2. For each path $Q \in \mathcal{J}$, let $x(Q)$ denote the last vertex of Q that lies on some alternating path starting in A_2; if no such vertex exists, let $x(Q)$ be the initial vertex of Q. We claim that

$$\mathcal{W} := \{\, Qx \mid Q \in \mathcal{J} \text{ and } x = x(Q) \,\}$$

is a wave in G; since $|A_2| > \kappa \geqslant 0$ and hence $\mathit{in}\,[\mathcal{W}] = \mathit{in}\,[\mathcal{J}] \subsetneqq A$, this wave $\mathcal{W}$ will be a hindrance and the lemma will be proved.

To show that $\mathcal{W}$ is a wave, we have to prove that $\mathit{ter}\,[\mathcal{W}]$ separates A from B. So let P be any A–B path. Since P is not an alternating path from A_2 to B_2, it meets $V[\mathcal{J}]$ and hence $V[\mathcal{W}]$; let y be its last vertex in $V[\mathcal{W}]$, and write $Q := Q_{\mathcal{J}}(y)$ and $x := x(Q)$. Suppose P avoids $\mathit{ter}\,[\mathcal{W}]$. Then $x \neq y$, and so there exists an alternating path R from A_2 to y which ends with an edge of $\mathcal{W}$. (Indeed, by definition of $\mathcal{W}$, there is an alternating A_2–x path R'; if x' is the first vertex of R' on $\mathring{y}Qx$, then $R'x'$ followed by yQx' in reverse order is an alternating path from A_2 to y.) By definition of $\mathcal{W}$, R avoids $V[\mathcal{J}]\backslash V[\mathcal{W}]$ and is thus an alternating path with respect to $\mathcal{W}$. Let z be the first vertex of R on yP. Then either $z = y$ or $z \notin V[\mathcal{W}]$, so RzP is again an alternating path with respect to $\mathcal{W}$. By definition of $\mathcal{W}$, RzP avoids $V[\mathcal{J}]\backslash V[\mathcal{W}]$. Therefore zP avoids $V[\mathcal{J}]\backslash\{ y \} \supseteq B_1$. (Recall that $y \notin B$, because $y \neq x$ and hence $y \in Q\mathring{x}$.) Thus RzP is an alternating path from A_2 to B_2, a contradiction.

We shall now assume that κ is infinite. To motivate our proof, let us consider the (much easier) case of $\mathcal{J} = \emptyset$. (This is an important special case, and we recommend that the reader remain aware of it throughout the proof of Lemma 1.) Assume Erdős's conjecture as true, and let S be an A-B separator as in the conjecture. Then $|S| \leqslant |B| = \kappa$. Let us think of the vertices in A as being 'to the left' of S, and of those in B as 'to its right'. Which other vertices of G will be to the left of S? Surely those which cannot be separated from A by $\leqslant \kappa$ vertices, i.e. which are joined to A by a fan of size $> \kappa$. We shall call these vertices 'popular'. If a popular vertex is in S, it is the starting vertex of a path to B that contains no other popular vertices; let us call such a path 'lonely', and its starting vertex 'special'. The special vertices, i.e. the vertices which are popular and from which we can get to B without hitting any other popular vertex, are in a sense 'rightmost' among the popular vertices, even when they are not in S. As we shall see, they turn out to be 'close enough' to S that they themselves form the set of endvertices of a hindrance in Γ, which is constructable without reference to S.

For the general case, we follow a similar approach, except that now all the relevant paths and fans will be alternating. Let us call a vertex $x \in V(G)\backslash A_1$ *popular* if either $x \in A_2$ or there exists an alternating A_2–x fan of order $> \kappa$. An alternating path P ending in B_2 and with no inner vertex in A_2 is called a *lonely path* if all its vertices are unpopular, *except* possibly its starting vertex and any vertices $x \in V[\mathcal{J}]$ such that, if e is the edge following x on P, then $e \notin E[\mathcal{J}]$. (In the latter case, the edge preceding x on P must be the edge of $\mathcal{J}$ starting in x.) Note that a final segment $x_i e_i x_{i+1} \ldots$ of a lonely path is again lonely if and only if x_i satisfies condition (iii) in the definition of an alternating path, i.e. if and only if $e_i \in E[\mathcal{J}]$ when $x_i \in V[\mathcal{J}]$.

Our first lemma is merely a technical argument that will be used twice later and has been extracted for economy. The first time we will use it is in the proof of Lemma 3.2 below, and for motivation the reader may prefer to read Lemma 3.2 and its proof first and then return to Lemma 3.1.

Lemma 3.1. *Let α be a cardinal, $\mathcal{L} = \{ L_\beta \mid \beta < \alpha \}$ a family of lonely paths, and $\mathcal{M} = \{ M_\beta \mid \beta < \alpha \}$ a family of pairwise disjoint alternating paths starting in A_2. Assume that, for each $\beta < \alpha$, the last vertex of M_β is the starting vertex of L_β. Then $\alpha \leqslant \kappa$.*

Proof. Suppose $\alpha > \kappa$. For each $\beta < \alpha$, let P_β be an alternating path from the starting vertex of M_β to the final vertex of L_β as provided by Lemma 2.2.

Construct an undirected forest $H = \bigcup_{\beta<\alpha} H_\beta$ from these paths, as follows. Let H_0 be the undirected graph underlying P_0. Now let $\beta < \alpha$ be given, and assume that H_γ has been defined for every $\gamma < \beta$. Let z_β denote the first vertex of P_β that is in $B_2 \cup V(H_\beta^-)$, where $H_\beta^- := \bigcup_{\gamma<\beta} H_\gamma$, and let H_β be the union of H_β^- with the undirected graph underlying $P_\beta z_\beta$. (If z_β occurs twice on P_β, we take $P_\beta z_\beta$ to stop at the first occurrence of z_β.) Since $|B_2| \leqslant \kappa$ and every path P_β ends in B_2, H has at most κ components. One of these components must have size $> \kappa$, so it contains a vertex z of degree $> \kappa$. Then z lies on $> \kappa$ of the paths P_β, so $z = z_\beta$ for every β in some set $\Delta \subseteq \alpha$ of size $> \kappa$.

Let

$$F := \{ P_\beta z \mid \beta \in \Delta \}.$$

Note that the paths in F are pairwise disjoint except for z, so F is an alternating A_2–z fan. Hence, z is popular. As the paths M_β are pairwise disjoint, we have $z \notin M_\beta$ for all but at most one $\beta \in \Delta$; let us delete this one β from Δ if it exists. Now for all $\beta \in \Delta$, we have that $z \in L_\beta$ and z is not the starting vertex of L_β (since this is on M_β).

Now consider any $\beta \in \Delta$. Since L_β is a lonely path and z is popular but not the starting vertex of L_β, we have $z \in V[\mathcal{J}]$, and if e denotes the edge following z on L_β then $e \notin E[\mathcal{J}]$. (Note that e exists, because $z \in V[\mathcal{J}]$ but L_β ends in B_2.) Since L_β is alternating, this means that the edge f preceding z on L_β must be the edge of $\mathcal{J}$ starting at z (and such an edge exists). Since $z \notin M_\beta$, the edge of P_β preceding z is precisely this edge f.

As β was chosen arbitrarily, this is true for every $\beta \in \Delta$ and thus contradicts the fact that for these β the paths $P_\beta \mathring{z}$ are disjoint. □

If the starting vertex of a lonely path is popular, then this vertex is called *special*; the set of all special vertices outside $V[\mathcal{J}]$ is denoted by S. Special vertices will be our prime candidates for the terminal vertices of the hindrance we are seeking to construct. Since the corresponding paths of the hindrance will have to be constructed from the fans connecting A_2 to these terminal vertices (making them popular), it is important that there are fewer special vertices to be connected in this way than there are connecting paths available from those fans.

Lemma 3.2. *There are at most κ special vertices.*

Proof. Suppose that $\{ s_\beta \mid \beta < \kappa^+ \}$ is a set of distinct special vertices, where κ^+ is the successor cardinal of κ. For each β, let L_β be a lonely path starting at s_β. Using the popularity of the s_β, we may inductively choose a family $\{ M_\beta \mid \beta < \kappa^+ \}$ of pairwise disjoint alternating paths M_β from A_2 to s_β. This contradicts Lemma 3.1. □

Let E denote the set of all those edges in G that lie on some lonely path, and let K be the graph

$$K := \bigcup \mathcal{J} - E.$$

Let $\mathcal{P}$ be the set of all those (undirected) components of K that contain a special vertex or a vertex from A. (Thus, $\mathcal{P}$ is a set of pairwise disjoint subpaths of paths in $\mathcal{J}$.) Let

$$T := \{ x \in V(G) \mid x \text{ is the last vertex on some } P \in \mathcal{P} \} .$$

We shall define our desired hindrance $\mathcal{W}$ in such a way that

$$ter[\mathcal{W}] = S \cup T.$$

Lemma 3.3. *If $P = x \dots y$ is a non-trivial component of K and $x \notin A$, then x is special (and hence $P \in \mathcal{P}$ and $y \in T$).*

Proof. Let r be the predecessor and s the successor of x on $Q(x)$. Then $rx \in E$, so there exists a lonely path starting at x with the edge rx. But preceding this path with s does not yield another lonely path (since $xs \in E(P)$, and hence $xs \notin E$). Therefore x must be popular (see the definition of lonely paths), and hence special. □

To construct $\mathcal{W}$, let us start from $\mathcal{P}$. Let $\mathcal{W}_0$ be the set of all paths $P \in \mathcal{P}$ that start in A. (These paths may be entire paths from $\mathcal{J}$, and they may be trivial.) Our aim is to complete $\mathcal{W}_0$ to our desired wave $\mathcal{W}$ by paths of the form $a \dots xPy$, where $a \in A_2$ and $P = x \dots y$ is a path as in Lemma 3.3, together with paths $a \dots s$ where again $a \in A_2$ and either $s \in S$ or s is a special vertex in $V[\mathcal{J}]$ making up a singleton component of K. It will not be possible to construct $\mathcal{W}$ in exactly this way, because the required paths may interfere with the paths in $\mathcal{W}_0$. However, such interference will be limited by Lemmas 3.1 and 3.2, and can therefore be overcome by the alternating path tools developed in Section 2.

Let

$$S' = \{ s_\zeta \mid \zeta < \nu \leqslant \kappa \}$$

be a well-ordering of those special vertices that are either in S or else are the initial vertex of some (possibly trivial) path $P \in \mathcal{P}$ (cf. Lemma 3.2). For each $\zeta < \nu$ in turn, we shall choose an alternating path P_ζ from A_2 to s_ζ, with the following properties:

(i) $P_\zeta \cap P_\xi = \emptyset$ for all $\xi < \zeta$;

(ii) $P_\zeta \cap Q_{\mathcal{J}}(s) \subseteq \{ s_\zeta \}$ for all $s \in S'$;

(iii) if $Q \in \mathcal{J}$ and $\xi < \zeta$ are such that $P_\xi \cap Q \neq \emptyset$, then $P_\zeta \cap Q \subseteq \{ s_\zeta \}$;

(iv) $E(P_\zeta) \cap E = \emptyset$.

Let $\zeta < \nu$ be given, and assume that paths P_ξ for all $\xi < \zeta$ have been chosen in accordance with (i)–(iv). By (ii), none of these paths contains $(s =)s_\zeta$. Since s_ζ is popular, there is an alternating A_2–s_ζ fan F of size $> \kappa$. Clearly, at most κ of the paths in F meet any of the paths P_ξ $(\xi < \zeta)$ or $Q(s)$ for $s \in S'$, except, for the latter, in s_ζ. Similarly, at most κ of the paths in F meet (in a vertex $\neq s_\zeta$) any path $Q \in \mathcal{J}$ that is hit by some P_ξ with $\xi < \zeta$. By Lemma 3.1, at most κ paths of F have an edge in E. (It is straightforward to check that the first edge in E on any path in F starts a lonely path.) We may thus choose P_ζ from the paths in F according to (i)–(iv).

Lemma 3.4. *For every $\zeta < \nu$, we have $E(P_\zeta) \cap E[\mathcal{J}] \subseteq E[\mathcal{W}_0]$. Thus, P_ζ is in fact an alternating path with respect to $\mathcal{W}_0$.*

Proof. If $e \in E(P_\zeta) \cap E[\mathcal{J}]$ then, by (iv) above, there is a component P of K containing e. By (ii), the initial vertex of P is not in S', and is therefore not special. By Lemma 3.3, therefore, the starting vertex of P must be in A, and so $P \in \mathcal{W}_0$. $\square$

Applying Lemma 2.1 ν times with the paths P_ζ, we now turn $\mathcal{W}_0$ into a warp from A onto *ter* $[\mathcal{W}_0] \cup S'$ with at most κ initial points in A_2. (Here we use that fact that, by (iii) above, no two of the alternating paths P_ζ use the same path in $\mathcal{W}_0$ to alternate on.) By (ii) above, the paths in $\mathcal{P}$ that start at the vertices in $S' \backslash S$ extend this warp to a warp $\mathcal{W}$. By Lemma 3.3, *ter* $[\mathcal{W}] = S \cup T$ as desired.

To prove that $\mathcal{W}$ is a wave in Γ, it remains to show that the set $S \cup T$ separates A from B; note that then $\mathcal{W}$ is also a hindrance, since $|$*in* $[\mathcal{W}] \cap A_2| = |S'| \leqslant \kappa$ by construction.

In order to prove that $S \cup T$ separates A from B, consider any A–B path $P = a \dots b$ in G. Suppose P avoids $S \cup T$.

Lemma 3.5. *Either $b \in B_2$, or b is the final vertex of an edge in $E \cap E[\mathcal{J}]$. In either case, b is not special but starts a lonely path.*

Proof. Suppose first that $b \in B_2$. Then b is not special, because $b \notin S$. Moreover, $\{ b \}$ is a trivial lonely path.

Suppose now that $b \in B_1$, and let $P' = x \dots b$ be the component of K containing b. As $b \notin T$, we have $P' \notin \mathcal{P}$, so $x \notin A$ and neither x nor b is special. By Lemma 3.3, therefore, P' is trivial, i.e. $b = x \notin A$. The edge e of $\mathcal{J}$ that ends in b is therefore in E, and hence lies on a lonely path. The final segment of this lonely path that starts at b (with e as its first edge) is again a lonely path, because $e \in E[\mathcal{J}]$. $\square$

Lemma 3.6. *The vertex a does not lie on a lonely path.*

Proof. If $a \in A_2$, then a is popular by definition, so being on—and hence starting—a lonely path would imply $a \in S$. If $a \in A_1$ and a lies on a lonely path, then this path uses the edge of $\mathcal{J}$ starting at a. Then $\{ a \}$ is a component of K, and hence a trivial path in $\mathcal{W}_0$ and in $\mathcal{W}$, giving $a \in T$. $\square$

Let x be the last vertex of P that is not on any lonely path, and let y be the vertex following x on P. Let L be a lonely path containing y. Then

(3.7) *$x \notin L$, and xyL is not a lonely path.*

Since $y \notin T$, y can only be in $V[\mathcal{J}]$ if the edge of $\mathcal{J}$ ending in y is in E. (Recall that L must use an edge of $\mathcal{J}$ incident with y, and apply Lemma 3.3.) We may therefore make the following assumption:

(3.8) *If $y \in V[\mathcal{J}]$, then L starts at y (with the edge of $\mathcal{J}$ that ends in y).*

Lemma 3.9. *The vertex y is popular.*

Proof. If y is not popular, then xyL can fail to be a lonely path only if it fails to be an alternating path. By (3.7) and (3.8), this can happen only if $x \in V[\mathcal{J}]$ and xyL fails to start with an edge of $\mathcal{J}$. But $x \notin B$, so x has a successor q on $Q(x)$. By (3.7) and (3.8), we have $q \neq y$. Now $qxyL$ is a lonely path (possibly containing q twice) that contradicts the choice of x. □

Let z be the last popular vertex on P. Then $z \neq b$, because b is unpopular by Lemma 3.5. As $z \in yP$ by Lemma 3.9, the choice of x and definition of y imply that z lies on some lonely path. But then $z \in V[\mathcal{J}]$, say $z \in Q \in \mathcal{J}$: otherwise the final segment of this lonely path that starts at z would again be lonely, and the popularity of z would mean that $z \in S$. Let q be the vertex following z on Q, and let t be the vertex following z on P.

Lemma 3.10. *$zq \notin E$.*

Proof. Let p be the vertex preceding z on Q. (This exists, since $z \neq a$.) If $zq \in E$, then $pz \notin E$: otherwise z would be not only popular but special, giving $\{ z \} \in \mathcal{P}$ and $z \in T$. But if $pz \notin E$, then $zq \in E$ implies by Lemma 3.3 that $z \in T$, a contradiction. □

Since $t \in yP$, there is a lonely path M containing t. As with y in (3.8), we may assume the following:

(3.11) *If $t \in V[\mathcal{J}]$, then M starts at t (with the edge of $\mathcal{J}$ that ends in t).*

By Lemma 3.10, zq is not an edge of M. By (3.11), this means that $t \neq q$; in particular, zq and zt are distinct edges. Moreover, zt is not an edge of M, since then M would have to use its starting edge again. Therefore, $qztM$ is an alternating path. Since t is not popular (by the choice of z), this means that $qztM$ is even a lonely path. (Note that zt is a 'real' edge, not the reverse of a $\mathcal{J}$-edge, so the popularity of z does not prevent this path from being lonely.) This, however, contradicts Lemma 3.10, completing the proof of Lemma 1.

4. Consequences

In this section we apply Lemma 1 to deduce some concrete partial results towards Erdős's conjecture. First, we need another lemma.

Lemma 4.1. *Let κ be an infinite cardinal. If Erdős's conjecture holds for all graphs of order $\leqslant \kappa$, then it holds for all webs $\Gamma = (G, A, B)$ such that $|A|, |B| \leqslant \kappa$.*

Proof. Let $\Gamma = (G, A, B)$ be a web with $|A|, |B| \leqslant \kappa$, and assume the conjecture holds for every graph of order $\leqslant \kappa$. Let G' be obtained from G by adding all edges xy such that G contains a set of $> \kappa$ independent x–y paths (i.e. paths that are disjoint except in x and y). To prove the conjecture for Γ, it suffices to find an orthogonal paths/separator pair $(\mathcal{P}, S)$ for $\Gamma' := (G', A, B)$. Indeed, then S is clearly also an A–B separator in G. As for the paths in $\mathcal{P}$, their foreign edges can be replaced inductively by paths in G whose interiors avoid each other and all the paths in $\mathcal{P}$ (since $|\mathcal{P}| \leqslant \kappa$), giving an A–B warp in G. We thus obtain an orthogonal pair for Γ.

Let G'' be the union of all minimal A–B paths in G'. (A path $P = a \dots b$ is *minimal* if G' contains no a–b path Q with $V(Q) \subsetneqq V(P)$.) It is now sufficient to find an orthogonal paths/separator pair for $\Gamma'' = (G'', A, B)$, which will clearly also be an orthogonal pair for Γ'. It thus suffices to show that $|G''| \leqslant \kappa$.

Suppose $|G''| > \kappa$, and consider a set $X \subseteq V(G'') \setminus (A \cup B)$ of size $> \kappa$, say $X = \{ x_\beta \mid \beta < \alpha \}$. (Recall that $|A|, |B| \leqslant \kappa$ by assumption.) For each $\beta < \alpha$, use the definition of G'' to find a minimal A–B path P_β in G' containing x_β. For all $\beta < \alpha$, define inductively P'_β as the maximal final segment of $P_\beta x_\beta$ that meets $\bigcup_{\gamma<\beta} P'_\gamma$ at most in its starting vertex s_β. Since $|A| \leqslant \kappa$, there is a vertex $s \in G''$ such that $s = s_\beta$ for every β in some set $\Delta \subseteq \alpha$ of size $> \kappa$. Then $F_1 := \bigcup_{\beta \in \Delta} s P_\beta x_\beta$ is a fan from s onto $Y := \{ x_\beta | \beta \in \Delta \}$.

Similarly, $\bigcup_{\beta \in \Delta} x_\beta P_\beta$ contains a fan F_2 from some set $Z \subseteq Y$ of size $> \kappa$ to a vertex t. Clearly, F_2 may be chosen so that no two of its paths meet a common path of F_1. It is then easy to combine F_1 and F_2 into a set of $> \kappa$ independent s–t paths in G''. Thus st is an edge of G', by definition of G'. But s and t are non-consecutive vertices on some common path P_β (take any β such that $x_\beta \in Z$), which contradicts the minimality of P_β. $\square$

Combining Lemma 1 and Lemma 4.1, we can now easily prove the following.

Theorem 4.2. *Let κ be an infinite cardinal. If Erdős's conjecture holds for all graphs of order $\leqslant \kappa$, then it holds for all webs $\Gamma = (G, A, B)$ such that $|A| \leqslant \kappa$.*

Proof. Let $\Gamma = (G, A, B)$ be a web with $|A| \leqslant \kappa$. Let $\overleftarrow{\mathcal{W}}$ be a wave in $\overleftarrow{\Gamma}$ such that $\overleftarrow{\Gamma'} := \overleftarrow{\Gamma} / \overleftarrow{\mathcal{W}}$ has no non-trivial wave. Let $B' := ter[\overleftarrow{\mathcal{W}}]$, and let $H \subseteq G$ be such that $\overleftarrow{\Gamma'} = (\overleftarrow{H}, B', A)$. (In other words, take the underlying graph of $\overleftarrow{\Gamma'}$, reverse its edges, and call the resulting graph H.) Since $\overleftarrow{\Gamma'}$ is unhindered, we have $|B'| \leqslant |A| \leqslant \kappa$ by Lemma 1. Now if the conjecture holds for all graphs of size $\leqslant \kappa$, then by Lemma 4.1 it holds for $\overleftarrow{\Gamma'}$, and there is a warp $\overleftarrow{\mathcal{J}}$ in $\overleftarrow{\Gamma'}$ together with a B'–A separator S in $\overleftarrow{H}$ consisting of a choice of one vertex from each path in $\overleftarrow{\mathcal{J}}$. But S is also a B–A separator in $\overleftarrow{G}$ (because $ter[\overleftarrow{\mathcal{W}}]$ is one) and hence an A–B separator in G. Thus S, together with $\mathcal{J}$ followed by a suitable subset of $\mathcal{W}$, is an orthogonal paths/separator pair for Γ. □

Corollary 4.3. *Erdős's conjecture is valid for all webs $\Gamma = (G, A, B)$ in which A is countable.*

Proof. By Theorem 4.2 and the fact that the conjecture holds for countable graphs [2]. □

Not surprisingly, Corollary 4.3 on its own does not need the full strength of Lemma 1. In fact, with hindsight, it is not too difficult to deduce the corollary directly from the main result of [3].

We conclude this section with an application of Lemma 1 to webs that come with a partial linkage.

Theorem 4.4. *Let $\Gamma = (G, A, B)$ be a web, and assume that G contains an A–B warp $\mathcal{J}$ such that $A \backslash in[\mathcal{J}]$ is countable. Then Erdős's conjecture holds for Γ.*

Proof. As in the proof of Theorem 4.2, we let $\overleftarrow{\mathcal{W}}$ be a wave in $\overleftarrow{\Gamma}$ such that $\overleftarrow{\Gamma'} := \overleftarrow{\Gamma}/\overleftarrow{\mathcal{W}}$ has no non-trivial wave. Let $B' := ter[\overleftarrow{\mathcal{W}}]$, and let $H \subseteq G$ be such that $\overleftarrow{\Gamma'} = (\overleftarrow{H}, B', A)$. Then $\overleftarrow{\Gamma'}$ is unhindered, and the final segments in $\overleftarrow{H}$ of the paths in $\overleftarrow{\mathcal{J}}$ form a B'–A warp $\overleftarrow{\mathcal{J}'}$ in $\overleftarrow{H}$. By Lemma 1, we have

$$|B' \backslash in[\overleftarrow{\mathcal{J}'}]| \leqslant |A \backslash ter[\overleftarrow{\mathcal{J}'}]| = |A \backslash in[\mathcal{J}]| \leqslant \aleph_0.$$

But such unhindered 'countable-like' webs as $\overleftarrow{\Gamma'}$ are linkable [3]. Let $\overleftarrow{\mathcal{L}}$ be a B'–A linkage in $\overleftarrow{H}$. The concatenations of the paths in $\mathcal{L}$ with their unique extensions in $\mathcal{W}$ then form an A–B warp in Γ, and B' is an A–B separator in G consisting of a choice of one vertex from each path in this warp. □

References

[[1]] R. Aharoni, König's duality theorem for infinite bipartite graphs, *J. London Math. Soc.* **29** (1984), 1–12.

[[2]] R. Aharoni, Menger's theorem for countable graphs, *J. Combin. Theory B* **43** (1987), 303–313.

[[3]] R. Aharoni, Linkability in countable-like webs, in: (G. Hahn et al., Eds.) *Cycles and Rays*, NATO ASI Ser. C, Kluwer Academic Publishers, Dordrecht 1990.

[[4]] R. Aharoni, Infinite matching theory, in: (R. Diestel, Ed.) *Directions in Infinite Graph Theory and Combinatorics*, Topics in Discrete Mathematics **3** (1992), North Holland.

[[5]] D. König, *Theorie der endlichen und unendlichen Graphen*, Akademische Verlagsgesellschaft, Leipzig 1936 (reprinted: Chelsea, New York 1950).

Ron Aharoni
Department of Mathematics
Technion
Haifa 32000
Israel

Reinhard Diestel
Mathematical Institute
Oxford University
Oxford OX1 3LB
England