

Proof of the Erdős-Menger conjecture with ends

Henning Bruhn Reinhard Diestel Maya Stein

Abstract

A well-known conjecture of Erdős states that, given an infinite graph G and sets $A, B \subseteq V(G)$, there exists a family of disjoint A - B paths \mathcal{P} together with an A - B separator X consisting of a choice of one vertex from each path in \mathcal{P} . There is a natural extension of this conjecture in which A , B and X may contain ends as well as vertices. We prove this extension by reducing it to the vertex version, which was recently proved by Aharoni and Berger.

1 Introduction

Erdős conjectured (see [6]) that Menger's theorem should extend to infinite graphs as follows:

Erdős-Menger Conjecture. *For every graph $G = (V, E)$ and any two sets $A, B \subseteq V$ there is a set \mathcal{P} of disjoint A - B paths in G and an A - B separator X consisting of a choice of one vertex from each of the paths in \mathcal{P} .*

A proof of this conjecture has recently been announced by Aharoni and Berger [1].

There is a natural extension of the Erdős-Menger conjecture in which the sets A and B may contain ends as well as vertices. Here, the A - B paths in \mathcal{P} can be either finite paths linking two vertices, or rays linking a vertex to an end, or double rays linking two ends. Similarly, the separator X may contain ends (that lie in A or B), thus blocking any ray belonging (= converging) to that end.

In this paper we prove the extended ends version of the conjecture by reducing it to the vertex version. Our proof uses a refinement of techniques developed in [3], where this reduction was carried out for countable graphs.

The Erdős-Menger conjecture for ends is not true for arbitrary sets A and B (of vertices and ends): a necessary condition is that the closure of A in the topological space $|G|$ associated with the graph G and its ends does not meet B , and vice versa. A comprehensive introduction to this space is given in [2], and we recall the basic definitions in Section 2. In those terms, our main result can be stated as follows:

Theorem 1.1. *Let $G = (V, E, \Omega)$ be a graph and let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, the closures being taken in $|G|$. Then G satisfies the Erdős-Menger conjecture for A and B .*

One may also consider a purely topological version of the Erdős-Menger conjecture, in which \mathcal{P} is any set of A - B arcs in the space $|G|$, and the set X is required to meet every A - B arc in $|G|$. This version of the conjecture can fail

unless A and B have disjoint closures in $|G|$. But in that case it reduces easily to Theorem 1.1; see [3].

2 Definitions and statement of result

The basic terminology we use can be found in [2]. All the graphs in this paper are simple and undirected. As most of the graphs we deal with are infinite, we recall a few (standard) concepts for infinite graphs. Let $G = (V, E)$ be a fixed infinite graph.

A 1-way infinite path is called a *ray*, a 2-way infinite path a *double ray*. The subrays of rays or double rays are their *tails*. The *ends* of G are the equivalence classes of rays under the following equivalence relation: two rays R_1, R_2 in G are *equivalent* if no finite set of vertices separates them. As one easily observes, this condition holds if and only if there are infinitely many disjoint (finite) R_1 – R_2 paths. This in turn is equivalent to the existence of a ray that meets both R_1 and R_2 infinitely often. The set of ends of G is denoted by $\Omega = \Omega(G)$, and we write $G = (V, E, \Omega)$ to refer to G together with its set of ends.

Paths in G can be finite paths (which contain at least one vertex), rays, double rays, or singleton sets $\{\omega\}$, where ω is an end of G .

We will now define the standard topology on G together with its ends. We start by viewing G itself as a 1-complex. (Thus, the basic open neighbourhoods of an inner point of an edge are the open intervals on the edge containing that point, while the basic open neighbourhoods of a vertex v are the unions of half-open intervals containing v , one from every edge at v .¹ The point set of this 1-complex will again be denoted by G .) To extend this topology to the set $G \cup \Omega$, we have to define a neighbourhood basis for every end $\omega \in \Omega$. To do so, consider any finite set $S \subseteq V$. Then $G - S$ has exactly one component $C = C(S, \omega)$ that contains a tail of every ray in ω . We say that ω *belongs* to C , and write $\overline{C}(S, \omega)$ for the component C together with all the ends of G belonging to C . As the basic open neighbourhoods of ω we now take all sets of the form

$$\hat{C}(S, \omega) := \overline{C}(S, \omega) \cup E'(S, \omega),$$

where S is any finite subset of V and $E'(S, \omega)$ is any union of half-edges $(z, y] \subseteq e$, one for every edge $e = xy$ with $x \in S$ and $y \in C$ (with $z \in \dot{e}$).

Let $|G|$ denote the topological space on $G \cup \Omega$ thus defined. (When G is locally finite, $|G|$ is a compact space known as the *Freudenthal compactification* of G ; see [2, 4, 5] for more.)

We write \overline{X} for the closure of a set $X \subseteq |G|$ in $|G|$. For example, the set $\overline{C}(S, \omega)$ defined above is the closure in $|G|$ of the set $C(S, \omega)$. Generally, the difference between a subgraph H and its closure \overline{H} is always a set of ends of G (possibly empty). These need not correspond to ends of H and should not be confused with them. For example, if G is the 1-way infinite ladder and H consists of all the rungs, then $\overline{H} \setminus H$ consists of one point, the unique end ω

¹Alternatively, we might fix for every edge $[u, v]$ a homeomorphism with the real interval $[0, 1]$ and take as basic open neighbourhoods for a vertex v only those unions of half-open edges $[v, z)$ whose images in $[0, 1]$ have the same length. This gives a different topology when vertices have infinite degrees, but since all relevant sequences of points will be sequences of vertices or of ends, the difference does not matter.

of G . But H itself has no ends. Similarly, the subgraph $H' = G - E(H)$ of G is a double ray and thus has two ends, but $\overline{H'} \setminus H' = \{\omega\}$ as before.

The union of a ray R and infinitely many disjoint paths starting on R but otherwise disjoint from R is a *comb* with *spine* R . The last points (vertices or ends) of those paths are the *teeth* of the comb. We will frequently use the following simple lemma:

Lemma 2.1. *In the graph $G = (V, E, \Omega)$ let R be a ray of an end ω , and let $X \subseteq V \cup \Omega$ such that $\omega \notin X$. Then $\omega \in \overline{X}$ if and only if G contains a comb with spine R and teeth in X . \square*

The closure of an infinite path P contains one or two ends of G . (Even if P is a double ray, its closure may contain only one end, as in the ladder example above.) We will often consider such an end as the first or last point of P , and when we say that two paths are *disjoint* then these points too shall be distinct. (The first and last point of a path $P = \{\omega\}$, of course, is ω .) For $A, B \subseteq V \cup \Omega$, a path is an A - B *path* if its first but no other point lies in A and its last but no other point lies in B .

A set $X \subseteq V \cup \Omega$ is an A - B *separator* in a subspace $T \subseteq |G|$ if every path P in T with its first point in A and its last point in B satisfies $\overline{P} \cap X \neq \emptyset$. (We express this informally by saying that “ P meets X ”, though strictly speaking we shall mean \overline{P} rather than just P .) We say that a set $Y \subseteq V \cup \Omega$ lies *on* a set \mathcal{P} of disjoint A - B paths if Y consists of a choice of exactly one vertex or end from every path in \mathcal{P} . We say that G *satisfies the Erdős-Menger conjecture for A and B* , or that *the Erdős-Menger conjecture holds for G, A, B* , if $|G|$ contains a set \mathcal{P} of disjoint A - B paths and an A - B separator on \mathcal{P} . (Thus, officially, we always refer to the ends version of the conjecture. But this is compatible with the traditional terminology: if neither A nor B contains an end then neither can any A - B path, so the conjecture with ends automatically defaults to the original conjecture in this case.)

The terms needed to state our main result are now precisely defined. We shall prove the following slight strengthening of Theorem 1.1 which, as in the vertex case, allows the intersection of A and B itself to be non-empty:

Theorem 2.2. *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap (\overline{B} \setminus B) = \emptyset = (\overline{A} \setminus A) \cap B$. Then G satisfies the Erdős-Menger conjecture for A and B .*

We remark that the disjointness condition in Theorem 2.2 is necessary [3]; a counterexample for when it is violated is given in [3]. The condition means that any ray whose end lies in A can be separated from B by a finite set of vertices, and vice versa with A and B interchanged. Note that this does not imply the much stronger condition that A and B can be finitely separated, in which case the proof is immediate by standard alternating path techniques [3]. A more typical example for the disjointness condition is to take as A and B distinct levels of vertices in a tree: if the tree is \aleph_0 -regular, for example, it contains infinitely many disjoint paths between these levels, so A and B have disjoint closures (in fact, are closed and disjoint) but cannot be finitely separated.

3 Proof of the theorem

Our aim is to reduce the ends version of the Erdős-Menger conjecture, Theorem 2.2, to the original vertex version as stated in the Introduction and recently proved by Aharoni and Berger.

We begin by showing that, as in the vertex case of the conjecture, we may assume without loss of generality that $A \cap B = \emptyset$. In the vertex case, one simply deletes $A \cap B$ from the graph, finds a path system and separator in $G - (A \cap B)$, and then adds the deleted vertices both to the path system (as singleton A - B paths) and to the separator, to obtain a solution for G . When $A \cap B$ is infinite, however, deleting it can result in the destruction or splitting of ends. Before we allow ourselves to assume that $A \cap B = \emptyset$, therefore, we have to make sure that this will not affect any ends in A or B . Our first lemma ensures this, and thereby reduces the stronger form of our theorem (Theorem 2.2) to the version stated in the introduction, Theorem 1.1.

Lemma 3.1. *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ satisfy*

$$A \cap (\overline{B} \setminus B) = \emptyset = (\overline{A} \setminus A) \cap B.$$

Then for the graph $G' := G - (A \cap B \cap V)$ there are sets $A', B' \subseteq V(G') \cup \Omega(G')$ satisfying the following conditions:

- (i) *if $A \subseteq V$ then $A' \subseteq A$, and if $B \subseteq V$ then $B' \subseteq B$;*
- (ii) *$A' \cap \overline{B'} = \emptyset = \overline{A'} \cap B'$;*
- (iii) *if G' satisfies the Erdős-Menger conjecture for A' and B' , then G satisfies it for A and B .*

Proof. Put $A' := A \setminus B$ and $B' := B \setminus A$, both of which are subsets of $|G|$. Consider a ray R of an end α in A' or B' , say in A' . Then R has a tail in G' . Indeed, if not then there are vertices of $A \cap B \cap V \subseteq B$ in every neighbourhood of $\alpha \in A \setminus B$. Consequently, $\alpha \in A \cap (\overline{B} \setminus B)$, which is a contradiction. Similarly, two rays R_1, R_2 in G' of which R_1 is a ray of an end $\omega \in A' \cup B'$ are equivalent in G' if and only if they are equivalent in G . Indeed, if R_1 and R_2 are equivalent in G then there is a ray $R_3 \in \omega$ that meets both of R_1 and R_2 infinitely often. Now R_3 has a tail in G' , showing that R_1 and R_2 are also equivalent in G' .

Thus, mapping every end of G in $A' \cup B'$ to the unique end of G' that contains tails of its rays defines a bijection between the ends in $A' \cup B'$ and certain ends in G' . Using this bijection (and a slight abuse of notation) we may view A' and B' also as subsets of $V(G') \cup \Omega(G')$. Clearly, these satisfy (i). Moreover, $A' \cap B'$ is still empty, so the disjointness assumption stated in the lemma implies (ii).

For (iii), let X' be an A' - B' separator on a set of disjoint A' - B' paths \mathcal{P}' in G' . Adding to \mathcal{P}' the trivial paths $\{x\}$ for all $x \in A \cap B$ yields a set \mathcal{P} of disjoint A - B paths with the A - B separator $X := X' \cup (A \cap B)$ on it. \square

In Lemma 3.5, we shall need a family of disjoint subgraphs of G (with certain properties) such that every end of A lies in the closure of one of these subgraphs. Such a family cannot always be found. But our next lemma finds instead a family of subgraphs such that the ends of A not contained in their closures form a set I that can be ignored: those ends will automatically be separated from B by any $(A \setminus I)$ - B separator on a set of disjoint A - B paths.

Lemma 3.2. *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Then for every set $A_\Omega \subseteq A \cap \Omega$ there exist a set $I \subseteq A_\Omega$, an ordinal μ^* , and families $(G_\mu)_{\mu < \mu^*}$ and $(S_\mu)_{\mu < \mu^*}$ such that, for every $\mu < \mu^*$, the graph $G_\mu - S_\mu$ is a component of $G - S_\mu$ with S_μ as its finite set of neighbours, and*

- (i) $\overline{G_\mu - S_\mu} \cap B = \emptyset$;
- (ii) if $G_\mu \neq \emptyset$ then $\overline{G_\mu} \cap A_\Omega \neq \emptyset$;
- (iii) $V(G_\nu \cap G_\mu) \subseteq S_\nu \cap S_\mu$ for all $\nu < \mu$.

Moreover,

- (iv) for every end $\alpha \in A_\Omega \setminus I$ there is a $\mu < \mu^*$ with $\alpha \in \overline{G_\mu}$;
- (v) every $(A \setminus I)$ - B separator on a set of disjoint $(A \setminus I)$ - B paths is also an A - B separator.

Proof. We construct the families $(G_\mu)_{\mu < \mu^*}$ and $(S_\mu)_{\mu < \mu^*}$ and a transfinite sequence $I_0 \subseteq I_1 \subseteq \dots \subseteq A_\Omega$ recursively. The sets I_μ ($\mu < \mu^*$) will serve as precursors to I . To simplify notation, we write $C_\mu := G_\mu - S_\mu$ for every μ . For the construction, we will in addition to (i)–(iii) require for every μ that

- (vi) $I_\mu \cap \overline{G_\nu} = \emptyset$ for all $\nu \leq \mu$.

We start by setting $I_0, G_0, S_0 := \emptyset$. Consider the least ordinal $\mu > 0$ such that the above sets are already defined for all $\lambda < \mu$. If μ is a limit, we set

$$I_\mu := \bigcup_{\lambda < \mu} I_\lambda$$

and $G_\mu, S_\mu := \emptyset$. This choice clearly satisfies (i)–(iii) and (vi).

Suppose now that μ is a successor, $\mu = \lambda + 1$ say. If every end in $A_\Omega \setminus I_\lambda$ lies in some $\overline{G_\nu}$ with $\nu < \mu$, we set $\mu^* := \mu$ and terminate the recursion. So suppose there is an end $\alpha \in A_\Omega \setminus I_\lambda$ that lies in no earlier $\overline{G_\nu}$. Then, if possible, choose a finite vertex set S such that $C(S, \alpha)$ avoids all G_ν with $\nu < \mu$.

Such a choice of S is impossible if and only if

$$\text{for every finite } S \subseteq V \text{ there is a } \nu < \mu \text{ with } C(S, \alpha) \cap G_\nu \neq \emptyset. \quad (1)$$

In this case we choose to ignore α , i.e. set $I_\mu := I_\lambda \cup \{\alpha\}$ and $G_\mu, S_\mu := \emptyset$. Again the requirements (i)–(iii) are clearly met, while (vi) holds by the choice of α .

Now suppose we can find S as desired. As $A \cap \overline{B} = \emptyset$, we can also find a basic open neighbourhood $\hat{C}(S', \alpha)$ of α in $|G|$ that is disjoint from B . We now define S_μ as the set of neighbours of $C(S \cup S', \alpha)$ and $G_\mu := G[S_\mu \cup C(S_\mu, \alpha)]$. Then (i) holds since $S_\mu \supseteq S'$, while (ii) holds as $\alpha \in \overline{G_\mu}$. To see (iii), first note that

$$G_\nu \cap C_\mu = \emptyset \text{ for all } \nu < \mu$$

by the choice of S . So, all we have to show is that $G_\nu \cap S_\mu \subseteq S_\nu$. Consider a vertex $v \in G_\nu \cap S_\mu$. Since S_μ is the set of neighbours of C_μ , there is a vertex $w \in C_\mu$ adjacent to v . As noted above, $w \notin G_\nu$. So v is a vertex in $G_\nu = C_\nu \cup N(C_\nu)$ with a neighbour outside G_ν , implying $v \notin C_\nu$ and hence $v \in S_\nu$, as desired.

Let us finally set $I_\mu := I_\lambda$ and verify (vi). We only need to show that $I_\mu \cap \overline{G_\mu} = \emptyset$. Suppose that intersection contains an end α' . Let $\mu' < \mu$ be minimal such that $\alpha' \in I_{\mu'}$. Then (1) should have been satisfied for μ' and α' , but fails with $S := S_\mu$ as $C(S_\mu, \alpha') = C_\mu$, a contradiction.

Having defined I_μ, G_μ and S_μ for all $\mu < \mu^*$ so that (i)–(iii) and (vi) are satisfied, we put

$$I := \bigcup_{\mu < \mu^*} I_\mu.$$

Together with the definition of μ^* this implies (iv). Observe that from (vi) we obtain $I \cap \overline{G_\mu} = \emptyset$ for all $\mu < \mu^*$.

To establish (v) let \mathcal{P} be a system of disjoint $(A \setminus I)$ – B paths and X an $(A \setminus I)$ – B separator on \mathcal{P} . Now suppose that X is not an A – B separator in $|G|$, i.e. there is a path Q from A to B that avoids X . By turning Q into a path \tilde{Q} from $A \setminus I$ to B that avoids X , we will obtain a contradiction.

We may assume that \tilde{Q} starts at an end $\alpha \in I$. Let μ be the step at which α was added to I , i.e. let μ be minimal with $\alpha \in I_\mu$. Choose a finite vertex set S such that $\overline{C}(S, \alpha)$ is disjoint from B (this is possible, as $A \cap \overline{B} = \emptyset$). Then any path of \mathcal{P} that meets $C(S, \alpha)$ must pass through S . Hence only finitely many paths of \mathcal{P} can meet $C(S, \alpha)$, and so $X_\alpha := X \cap \overline{C}(S, \alpha)$ is also finite. Conditions (iii) and (iv) ensure that every end in X_α lies in exactly one \overline{C}_λ ; let $\{\lambda_1, \dots, \lambda_m\}$ be the set of these λ . Then for

$$S' := S \cup (X_\alpha \cap V) \cup \bigcup_{i=0}^m S_{\lambda_i}$$

we have

$$\overline{C}(S', \alpha) \cap X = \emptyset.$$

Now, all we need is a point of $A \setminus I$ that lies in $\overline{C}(S', \alpha)$ (and thus can be used to change Q into the desired path). Indeed, if there is an ordinal $\lambda < \mu$ such that $G_\lambda \neq \emptyset$ and

$$C_\lambda \subseteq C(S', \alpha), \tag{2}$$

we can complete the proof as follows. By (ii) for λ there will be an end $\alpha' \in A$ in $\overline{C}_\lambda \subseteq \overline{C}(S', \alpha)$. Since $I \cap \overline{G}_\lambda = \emptyset$, we have $\alpha' \in A \setminus I$. Take an α' – Q path P in $\overline{C}(S', \alpha)$ with last vertex x , say. Then P avoids X , and hence so does the path $\tilde{Q} := PxQ$. Thus, \tilde{Q} is as desired.

So suppose there is no ordinal $\lambda < \mu$ satisfying (2). Then for all $\lambda < \mu$ we have either $C_\lambda \cap C(S', \alpha) = \emptyset$ or $C_\lambda \cap S' \neq \emptyset$. As all the C_λ are disjoint by (iii), only finitely many of them meet S' ; let $\lambda_{m+1}, \dots, \lambda_n$ be the corresponding ordinals. Then

$$S'' := S' \cup \bigcup_{i=m+1}^n S_{\lambda_i}$$

satisfies $C(S'', \alpha) \cap C_\lambda = \emptyset$ for all $\lambda < \mu$.

However, $G_\lambda \cap C(S'', \alpha)$ cannot be empty for all $\lambda < \mu$, as this would contradict (1) for step μ with $S := S''$. So there exists an ordinal $\lambda < \mu$ with $S_\lambda \cap C(S'', \alpha) \neq \emptyset$. A vertex v in this intersection must have a neighbour in C_λ , which then also lies in $S' \cup C(S', \alpha)$ because $C(S'', \alpha) \subseteq C(S', \alpha)$. Thus,

$$(S' \cup C(S', \alpha)) \cap C_\lambda \neq \emptyset.$$

Since $C_\lambda \not\subseteq C(S', \alpha)$ by assumption, this implies that C_λ meets S' . But then $\lambda \in \{\lambda_{m+1}, \dots, \lambda_n\}$ and hence $S_\lambda \subseteq S''$, contradicting the fact that v lies in both S_λ and $C(S'', \alpha)$. \square

For our end-to-vertex reduction we need two more lemmas.

Lemma 3.3. [3] *Let H be a subgraph of a graph G , let $S \subseteq V(H)$ be finite, and let $T \subseteq V(H) \cup \Omega(G)$ be such that $T \subseteq \overline{H}$. Then \overline{H} contains a set \mathcal{P} of disjoint S - T -paths and an S - T -separator (in \overline{H}) on \mathcal{P} .*

For a set T of vertices in a graph H , a T -path is a path that meets T only in its first and last vertex. A set of paths will be called *disjoint outside* a given subgraph $Q \subseteq H$ if distinct paths meet only in Q .

Lemma 3.4. [7, 3] *Let H be a graph, $T \subseteq V(H)$ finite, and $k \in \mathbb{N}$. Then H has a subgraph H' containing T such that for every T -path $Q = s \dots t$ in H meeting $H - H'$ there are k distinct T -paths from s to t in H' that are disjoint outside Q .*

Our next lemma allows us to replace the set $A \subseteq V \cup \Omega$ in Theorem 1.1 with a set A' consisting only of vertices.

Lemma 3.5. *Let $G = (V, E, \Omega)$ be a graph, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Then there exist a minor $G' = (V', E', \Omega')$ of G and sets $A' \subseteq V'$ and $B' \subseteq V' \cup \Omega'$ satisfying the following conditions:*

- (i) *if $B \subseteq V$ then $B' \subseteq B$;*
- (ii) *$A' \cap \overline{B'} = \emptyset = \overline{A'} \cap B'$;*
- (iii) *G satisfies the Erdős-Menger-conjecture for A and B if G' satisfies it for A' and B' .*

Proof. Applying Lemma 3.2 with $A_\Omega := A \cap \Omega$ we obtain an ordinal μ^* , subgraphs G_μ , finite vertex sets S_μ and a set of ends $I \subseteq A$. Our aim is to change G into G' by deleting and contracting certain connected subgraphs of our graphs $G_\mu - S_\mu$. By Lemma 3.2 (iii) we shall be able to do this independently for the various G_μ : for each $\mu < \mu^*$ separately, we shall find in $G_\mu - S_\mu$ a set $\mathcal{D}_1(\mu)$ of connected subgraphs to be deleted, and another set $\mathcal{D}_2(\mu)$ of connected subgraphs that will be contracted.

Fix $\mu < \mu^*$. If G_μ is empty we let $\mathcal{D}_1(\mu) = \mathcal{D}_2(\mu) = \emptyset$. Assume now that $G_\mu \neq \emptyset$. Put $A_\mu := A \cap \overline{G_\mu}$. Applying Lemma 3.3 to $H = G_\mu$ we find in $\overline{G_\mu}$ a finite set \mathcal{P} of disjoint S_μ - A_μ paths and an S_μ - A_μ separator X_μ on \mathcal{P} . We write $X_\mu = U_\mu \cup O_\mu$, where $U_\mu = X_\mu \cap V$ and $O_\mu = X_\mu \cap \Omega$, both of which are finite since $|X_\mu| \leq |\mathcal{P}| \leq |S_\mu|$. Moreover,

$$U_\mu \text{ separates } S_\mu \text{ from } A_\mu \setminus O_\mu \text{ in } G. \quad (3)$$

Indeed, every S_μ - $(A_\mu \setminus O_\mu)$ path in G lies in $\overline{G_\mu}$ and hence meets X_μ , and since it cannot meet O_μ unless it ends there, it meets X_μ in U_μ .

We define $\mathcal{D}_1(\mu)$ as the set of all the components D of $G - U_\mu$ whose closure \overline{D} meets $A_\mu \setminus O_\mu$. By (3), these components satisfy $D \subseteq G_\mu - S_\mu$, and their neighbourhood $N(D) \subseteq U_\mu$ in G is finite. In addition,

$$\overline{D} \cap O_\mu = \emptyset \text{ for all } D \in \mathcal{D}_1(\mu). \quad (4)$$

For if $\alpha \in \overline{D} \cap O_\mu$, say, and P is the S_μ - A_μ path in \mathcal{P} that ends in α , then P has a tail in D . Since P does not meet $U_\mu \supseteq N(D)$, this implies $P \subseteq \overline{D}$. Consequently, $S_\mu \cap D$ is not empty as it contains at least the first vertex of P . This contradicts $D \subseteq G_\mu - S_\mu$.

Put

$$H_\mu := G_\mu - \bigcup \mathcal{D}_1(\mu).$$

Note that, as every $v \in U_\mu$ lies on a path in \mathcal{P} ,

G_μ contains a set of disjoint H_μ - A_μ paths whose set of first points is U_μ . (5)

By (3) and the definition of H_μ , we have $\overline{H_\mu} \cap A \subseteq U_\mu \cup O_\mu = X_\mu$. Since O_μ is finite, we can extend $U_\mu \cup S_\mu$ to a finite set $T_\mu \subseteq V(H_\mu)$ that separates the ends in O_μ pairwise in G . Let H'_μ be the finite subgraph of H_μ containing T_μ which Lemma 3.4 provides for $k := |S_\mu| + 1$, and for each $\alpha \in O_\mu$ let D_α be the component of $G - H'_\mu$ to which α belongs. Finally, we conclude our definitions for μ by setting $\mathcal{D}_2(\mu) := \{D_\alpha \mid \alpha \in O_\mu\}$.

Define for $i = 1, 2$

$$\mathcal{D}_i := \bigcup_{\mu < \mu^*} \mathcal{D}_i(\mu).$$

Observe that, by Lemma 3.2 (iii) and since their neighbourhoods in G are finite, the elements of $\mathcal{D}_1 \cup \mathcal{D}_2$ have pairwise disjoint closures.

Before we can define G' , we first have to introduce a graph $\tilde{G} = (\tilde{V}, \tilde{E}, \tilde{\Omega})$ from which we will obtain G' by deleting certain vertices. Let \tilde{G} be obtained from $G - \bigcup \mathcal{D}_1$ by contracting every $D_\alpha \in \mathcal{D}_2$ to a single vertex a_α , and put

$$A^* := \{a_\alpha \mid D_\alpha \in \mathcal{D}_2\}.$$

Then for $Z := \bigcup \mathcal{D}_1 \cup \bigcup \mathcal{D}_2$ we have

$$G - Z = G \cap \tilde{G} = \tilde{G} - A^*.$$

By Lemma 3.2 (iii) and by (3), the union of the sets of paths in (5) for all $\mu < \mu^*$ is a set of disjoint paths. Thus, for $U := \bigcup_{\mu < \mu^*} U_\mu$

there is a set of disjoint U - A paths whose set of first points is U , and whose paths meet \tilde{G} only in U . (6)

An important property of \tilde{G} is that the ends of G in $B \cap \Omega$ correspond closely to ends of \tilde{G} . To establish this correspondence formally, we begin with the following observation:

Every ray of an end $\beta \in B$ has a tail in $G - Z$. (7)

To see this, recall that all the $D \in \mathcal{D}_1 \cup \mathcal{D}_2$ have pairwise disjoint closures, and that each of them is a connected subgraph of G whose closure contains an end or a vertex of A . Hence, a ray R of β meets only finitely many $D \in \mathcal{D}_1 \cup \mathcal{D}_2$, as we could otherwise find infinitely many disjoint R - A paths, giving $\overline{A} \cap B \neq \emptyset$ by Lemma 2.1 – a contradiction. Also, R meets every $D \in \mathcal{D}_1 \cup \mathcal{D}_2$ only finitely often. Indeed, D lies in G_μ for some $\mu < \mu^*$ and is thus, by Lemma 3.2 (i), separated from β by its finite set of neighbours $N(D)$. This establishes (7).

Let R_1, R_2 be two rays in $G \cap \tilde{G}$, and assume that the end of R_1 lies in B . Then R_1 and R_2 are equivalent in G if and only if they are equivalent in \tilde{G} . (8)

To prove (8), suppose first that R_1, R_2 are equivalent in G , i.e. belong to the same end $\beta \in B$. Then there is a ray R_3 that meets both R_1 and R_2 infinitely often, and hence ends in β . By (7), R_3 has a tail in $G - Z = \tilde{G} - A^*$, showing that R_1 and R_2 are equivalent also in \tilde{G} .

Conversely, if R_1 and R_2 are joined in \tilde{G} by infinitely many disjoint paths, we can replace any vertices $a_\alpha \in \tilde{V} \setminus V = A^*$ on these paths by finite paths in D_α to obtain infinitely many disjoint R_1 - R_2 paths in G . This completes the proof of (8).

We can now define our correspondence between the ends in B and certain ends of \tilde{G} . For every end $\beta \in B$ there is by (7) an end $\beta' \in \tilde{\Omega}$ such that $\beta \cap \beta' \neq \emptyset$. By (8), this end β' is unique and the map $\beta \mapsto \beta'$ is injective. Moreover,

$$\tilde{B} := (B \cap V) \cup \{\beta' \mid \beta \in B \cap \Omega\} \subseteq \tilde{V} \cup \tilde{\Omega}$$

by Lemma 3.2 (i). For each $\mu < \mu^*$, let

$$\tilde{A}_\mu := U_\mu \cup \{a_\alpha \mid \alpha \in O_\mu\},$$

if $G_\mu \neq \emptyset$; if $G_\mu = \emptyset$, put $A_\mu, \tilde{A}_\mu := \emptyset$. Then let

$$\tilde{A} := \left(A \setminus \left(\bigcup_{\mu < \mu^*} A_\mu \cup I \right) \right) \cup \bigcup_{\mu < \mu^*} \tilde{A}_\mu,$$

which is a subset of \tilde{V} by Lemma 3.2 (iii),(iv). Finally, let

$$G' := \tilde{G} - (\tilde{A} \cap \tilde{B}).$$

To show the assertions (i)–(iii), we will apply Lemma 3.1 to the graph \tilde{G} and the sets \tilde{A} and \tilde{B} .

So, let us show that

$$(\overline{\tilde{A}} \setminus \tilde{A}) \cap \tilde{B} = \emptyset = \tilde{A} \cap (\overline{\tilde{B}} \setminus \tilde{B})$$

(with closures taken in $|\tilde{G}|$). We trivially have $\tilde{A} \cap (\overline{\tilde{B}} \setminus \tilde{B}) = \emptyset$ because $\tilde{A} \subseteq \tilde{V}$. To prove that $(\overline{\tilde{A}} \setminus \tilde{A}) \cap \tilde{B} = \emptyset$, consider an end $\beta' \in \tilde{B}$. The corresponding end $\beta \in B$ has a neighbourhood $C := \hat{C}(S, \beta)$ in $|G|$ that avoids A . By (6), and since S is finite, the intersection $C \cap U =: U_C$ is finite. Also, as in the proof of (7), C may meet only finitely many $D_\alpha \in \mathcal{D}_2$. Denote by O_C the set of the corresponding $a_\alpha \in \tilde{G}$. Adding to $S \setminus Z$ the sets U_C and O_C then yields a finite set $S' \subseteq \tilde{V}$ such that the neighbourhood $\hat{C}'(S', \beta')$ in $|\tilde{G}|$ even avoids \tilde{A} .

Thus, Lemma 3.1 is applicable and yields sets $A' \subseteq V'$ and $B' \subseteq V' \cup \Omega'$ satisfying (ii). Assertion (i) follows from the definition of \tilde{B} and Lemma 3.1 (i).

We now prove assertion (iii) of the lemma. Suppose G' satisfies the Erdős-Menger conjecture for A' and B' . Then, by Lemma 3.1, there is also in \tilde{G} a set $\tilde{\mathcal{P}}$ of disjoint \tilde{A} - \tilde{B} paths and an \tilde{A} - \tilde{B} separator \tilde{X} on $\tilde{\mathcal{P}}$. In order to turn $\tilde{\mathcal{P}}$

into a set $\mathcal{P} := \{P \mid \tilde{P} \in \tilde{\mathcal{P}}\}$ of disjoint A - B paths in G , consider any $\tilde{P} \in \tilde{\mathcal{P}}$. If the first point a of \tilde{P} lies in A we leave \tilde{P} unchanged, i.e. set $P := \tilde{P}$. If $a \in \tilde{A} \setminus (A \cup A^*)$, then $a \in U_\mu$ for some $\mu < \mu^*$, and we let P be the union of \tilde{P} with an A_μ - U_μ path in G_μ that ends in a ; this can be done disjointly for different $\tilde{P} \in \tilde{\mathcal{P}}$ if we use the paths from (6). Moreover, the A_μ - H_μ path concatenated with \tilde{P} in this way has only its last vertex in \tilde{G} , so it will not meet any other vertices on $\tilde{\mathcal{P}}$. Finally if $a = a_\alpha \in A^*$, we let P be obtained from \tilde{P} by replacing a with a path in D_α that starts at the end α and ends at the vertex of D_α incident with the first edge of \tilde{P} (the edge incident with a). In all these cases we have $P \subseteq G$, because \tilde{P} has no vertex in A^* other than possibly a . And no vertex of P other than possibly its last vertex lies in B , because $B \cap V = \tilde{B} \cap \tilde{V}$ and any new initial segment of P lies in a subgraph $G_\lambda - S_\lambda$ of G which avoids B by Lemma 3.2 (i).

It remains to check that the paths P just defined have distinct last points in B even when the last points of the corresponding paths \tilde{P} are ends. However if \tilde{P} ends in $\beta' \in \tilde{B}$ then its tail $\tilde{P} - a \subseteq P \subseteq G$ is equivalent in \tilde{G} to some ray in $\beta' \cap \beta$, by definition of β' . By (8) this implies $\tilde{P} - a \in \beta$, so the last point of P is $\beta \in B$. And since the map $\beta \mapsto \beta'$ is well defined, these last points differ for distinct P , because the corresponding paths \tilde{P} have different endpoints β' by assumption.

We still need an A - B separator on \mathcal{P} . The only vertices $x \in \tilde{X}$ that do not lie on the path P obtained from the path \tilde{P} containing x are points in A^* . So let X be obtained from \tilde{X} by replacing every end $\beta' \in \tilde{X} \cap \tilde{B}$ with the corresponding end $\beta \in B$ and replacing every $a_\alpha \in \tilde{X} \cap A^*$ with the end $\alpha \in A$. Since $P \in \mathcal{P}$ starts in α if \tilde{P} starts in a_α (and P ends in β if \tilde{P} ends in β'), this set X consists of a choice of one point from every path in \mathcal{P} .

Let us then show that

$$X \text{ is an } A\text{-}B \text{ separator in } G. \quad (9)$$

Suppose there exists a path $Q \subseteq G - X$ that starts in A and ends in B . Lemma 3.2 (v) enables us to choose Q as a path starting in $A \setminus I$. Our aim is to turn Q into an \tilde{A} - \tilde{B} path Q' in \tilde{G} that avoids \tilde{X} , which contradicts the choice of \tilde{X} .

If Q meets $\bigcup \mathcal{D}_1$, it has a last vertex there by (7), in $D \in \mathcal{D}_1(\lambda)$, say. Its next vertex a lies in U_λ , by the definition of D . We then define (for the time being) Q' as the final segment aQ of Q starting at a . If Q has no vertex in $\bigcup \mathcal{D}_1$, then either the first point of Q is a vertex $a \in A \cap \tilde{A}$ (in which case we put $Q' := Q$), or Q starts at an end $\alpha \in A \setminus I$. By Lemma 3.2 (iv), there exists a $\lambda < \mu^*$ such that $\alpha \in \overline{G}_\lambda$, which implies $\alpha \in O_\lambda$. We make $a := a_\alpha$ the starting vertex of Q' and continue Q' along Q , beginning with the last D_α - \tilde{G} edge on Q . Our assumption of $\alpha \notin X$ implies that $a_\alpha \notin \tilde{X}$, by the definition of X . Thus in the first two cases, Q' is now a path in $G - \bigcup \mathcal{D}_1$; in the third, Q' is a path in $(G - \bigcup \mathcal{D}_1)/D_\alpha$, which starts at the vertex $a \in \tilde{A}$ and avoids \tilde{X} .

However, Q' may still meet \mathcal{D}_2 . And although we know from (7) that Q' has a last vertex in $\bigcup \mathcal{D}_2$, say in $D_{\alpha'}$, we cannot simply shorten Q' to a path $a_{\alpha'}Q'$ in \tilde{G} , because it may happen that $a_{\alpha'} \in \tilde{X}$. Instead, we will use Lemma 3.4 to replace any segments of Q' that meet some $D_\alpha \in \mathcal{D}_2$ (with $a_\alpha \neq a$) by paths through the corresponding G_μ that avoid \tilde{X} . As we only have to deal with a finite initial segment of Q' and the D_α are all disjoint, we are able to modify Q'

step by step. Eventually, we will obtain a (walk that can be pruned to a) path Q' in \tilde{G} that avoids \tilde{X} , yielding the desired contradiction.

So consider a segment of Q' that meets some $D_\alpha \in \mathcal{D}_2$. By definition of D_α we may assume that segment to be a T_μ -path $sQ't$ in H'_μ , where μ is such that $D_\alpha \subseteq G_\mu$. By definition of H'_μ (which is a subgraph of \tilde{G} by Lemma 3.2 (iii), i.e. no parts of H'_μ were deleted or contracted when we defined \tilde{G}), there are $|S_\mu| + 1$ paths from s to t in H'_μ that are disjoint outside $sQ't$. But H'_μ contains at most $|S_\mu|$ vertices from \tilde{X} : since these lie on disjoint paths ending in \tilde{B} and S_μ separates $H'_\mu \subseteq G_\mu$ from B in G and hence from \tilde{B} in \tilde{G} , all of these paths must meet S_μ . So one of our $|S_\mu| + 1$ s - t paths in H'_μ avoids \tilde{X} , and we can use this path to replace $sQ't$ on Q' . This completes the proof of (9). \square

Proof of Theorem 2.2: Let $G = (V, E, \Omega)$ be given, and let $A, B \subseteq V \cup \Omega$ be such that $A \cap (\overline{B} \setminus B) = \emptyset = (\overline{A} \setminus A) \cap B$. By Lemma 3.1, we may assume that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Applying Lemma 3.5 twice, first for A and then for B , we may further assume that $A \cup B \subseteq V$. Now the statement to be proved is Erdős's conjecture as stated in the Introduction, which has been proved by Aharoni and Berger [1]. \square

References

- [1] R. Aharoni and E. Berger. Proof of the Erdős-Menger conjecture. Communication at the Oberwolfach meeting on *Graph Theory*, 2005.
- [2] R. Diestel. *Graph Theory* (3rd edition). Springer-Verlag 2005 (to appear).
- [3] R. Diestel. The countable Erdős-Menger conjecture with ends. *J. Combin. Theory (Series B)*, 87:145–161, 2003.
- [4] R. Diestel. The cycle space of an infinite graph. To appear in *Comb., Probab., Computing*.
- [5] R. Diestel and D. Kühn. Graph-theoretical versus topological ends of graphs. *J. Combin. Theory (Series B)*, 87:197–206, 2003.
- [6] C.St.J.A. Nash-Williams, Infinite graphs, a survey, *J. Combin. Theory (Series B)*, 3:286–301, 1967.
- [7] M. Stein. Die Erdős-Menger'sche Vermutung für abzählbare Graphen mit Enden. Diplomarbeit Universität Hamburg 2002.

Version 22 January, 2005

Henning Bruhn <hbruhn@nullsinn.net>
 Reinhard Diestel <diestel@math.uni-hamburg.de>
 Maya Stein <fm7y052@public.uni-hamburg.de>

Mathematisches Seminar, Universität Hamburg, Bundesstraße 55, 20146 Hamburg, Germany