

The countable Erdős-Menger conjecture with ends

Reinhard Diestel

Erdős conjectured that, given an infinite graph G and vertex sets $A, B \subseteq V(G)$, there exist a set \mathcal{P} of disjoint A - B paths in G and an A - B separator X ‘on’ \mathcal{P} , in the sense that X consists of a choice of one vertex from each path in \mathcal{P} . We prove, for countable graphs G , the extension of this conjecture in which A, B and X are allowed to contain ends as well as vertices, and where the closure of A avoids B and vice versa. (Without the closure condition the extended conjecture is false.)

1. Introduction

Most graph theorists with an interest in infinite graphs will, I expect, follow C.St.J.A. Nash-Williams in thinking of the following conjecture of Erdős as the main open problem in infinite graph theory:

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} .*

The conjecture appears in print first in Nash-Williams’s 1967 survey [14] on infinite graphs, although it seems to be considerably older. It was proved by Aharoni [2] for countable graphs, by Aharoni, Nash-Williams and Shelah [1, 5] for bipartite graphs with bipartition (A, B) , and in [4] for arbitrary graphs with A countable. The current state of the art, including further partial results, is described in Aharoni [3].

In this paper we consider the natural extension of the conjecture to *ends*, and prove the extension for countable G . (Note that G may have uncountably many ends.) Thus, A and B will be allowed to contain ends as well as vertices, the A - B paths in \mathcal{P} may be rays or double rays joining the appropriate vertices or ends in A and B , and the separator X may likewise contain ends from $A \cup B$.

As the precise statement of our theorem relies on the terms to be introduced in Section 2, we defer it to Section 3. We shall also derive some natural corollaries in Section 3, and discuss related results and open problems. The proof of our main result is given in Sections 4 and 5.

2. Terminology

The basic terminology we use is that of [7] – except that most of our graphs will be infinite, and $|G|$ will denote a certain topological space associated with a graph G , not its order. Our graphs are simple and undirected, although the result we prove can easily be adapted to directed graphs. Let $G = (V, E)$ be a countable graph, fixed throughout the paper.

We think of *paths* in G as subgraphs together with the natural linear ordering of their vertices. Thus, an infinite path $x_0x_1x_2\dots$ that has a first but no last vertex differs from the path $\dots x_2x_1x_0$ that has a last but no first vertex, although both these paths are identical as graphs. Paths of the former type are called *rays*; paths of the latter type are *inverse rays*. Paths with neither a first nor a last vertex are *double rays*, and the subrays of rays or inverse rays or double rays are their *tails*. In addition to the above, we also call singleton sets $\{\omega\}$, where ω is an end of G (defined below), a *path* in G .

Two rays in G are *equivalent* if no finite set of vertices separates them in G . The corresponding equivalence classes of rays are the *ends* of G ; the set of these ends is denoted by $\Omega = \Omega(G)$, and G together with its ends is referred to as $G = (V, E, \Omega)$. (The grid, for example, has one end, the double ladder has two, and the binary tree has continuum many; see [10] for more background.) We shall endow our graph G , complete with vertices, edges and ends, with a (now standard) topology first introduced by Jung [13], to be defined below. This topological space will be denoted by $|G|$, and the closure in $|G|$ of a subset X will be written as \overline{X} .

To define $|G|$, we start with G viewed as a 1-complex. Then every edge is homeomorphic to the real interval $[0, 1]$, the basic open sets around an inner point being just the open intervals on the edge. The basic open neighbourhoods of a vertex x are the unions of half-open intervals $[x, z]$, one from every edge $[x, y]$ at x ; note that we do not require local finiteness here.

For $\omega \in \Omega$ and any finite set $S \subseteq V$, the graph $G - S$ has exactly one component $C = C(S, \omega)$ that contains a tail of every ray in ω . We say that ω *belongs* to C . Write $\Omega(S, \omega)$ for the set of all ends of G belonging to C , and $E(S, \omega)$ for the set of all edges of G between S and C . Now let $|G|$ be the point set $V \cup \Omega \cup \bigcup E$ endowed with the topology generated by the open sets of the 1-complex G and all sets of the form

$$\widehat{C}(S, \omega) := C(S, \omega) \cup \Omega(S, \omega) \cup E'(S, \omega),$$

where $E'(S, \omega)$ is any union of half-edges $(x, y] \subset e$, one for every $e \in E(S, \omega)$, with $x \in \dot{e}$ and $y \in C$. (So for each end ω , the sets $\widehat{C}(S, \omega)$ with S varying over the finite subsets of V are the basic open neighbourhoods of ω .) This is the standard topology on graphs with ends, denoted in [8] as TOP. With this topology, $|G|$ is a Hausdorff space in which every ray converges to the end that contains it. $|G|$ is easily seen to be compact if and only if every vertex has finite degree.

A *subgraph* $H = (U, F)$ of G will be viewed topologically as just the point set $U \cup \bigcup F$, without any ends. Then the closure \overline{H} of this set in $|G|$ may contain some ends of G , but these should not be confused with ends of H . (In fact, the only ends we shall ever consider seriously in this paper will be ends of G . From a graph-theoretic viewpoint this may seem unorthodox when considering subgraphs, but it is natural topologically and simplifies matters greatly.)

A *subspace* of $|G|$ is any subset of the point set $|G|$ with the subspace topology. The only subspaces we shall consider, however, will be subgraphs of G , possibly together with some ends. That is to say, if a subspace we consider contains an inner point of an edge e , it will contain the entire edge e together with its endvertices.

When we speak about paths informally, we shall by default mean their closures in G . Thus, every path P has two endpoints (a *first* and a *last* point, which are formally the two boundary points of \overline{P} and may be either vertices or ends), and these are taken into account when we speak about disjointness, containment etc. Thus, *disjoint* paths must have distinct endpoints, a path *in* a subspace T of $|G|$ must have its endpoints in T , and so on. 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 . Note that a path starting in an end $\alpha \in A$ and ending in B need not have a last point in A , in which case it will not contain an *A–B* path (and similarly for paths ending in an end $\beta \in B$).

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 last point in B (which need not be or contain an *A–B* path) has a point in X , ie. $\overline{P} \cap X \neq \emptyset$. We say that X lies *on* a set \mathcal{P} of disjoint *A–B* paths if X consists of a choice of exactly one point from every path in \mathcal{P} .

The following lemma will be used repeatedly in our proofs:

Lemma 2.1. *If $R \subseteq G$ is a ray and $X \subseteq |G|$ a set such that G contains infinitely many disjoint paths starting on R and ending in X , then the end ω of R lies in \overline{X} , the closure of X in $|G|$.*

Proof. If $\omega \notin \overline{X}$, then ω has a neighbourhood $\widehat{C}(S, \omega)$ in $|G|$ that avoids X . As $R \in \omega$, R has a tail in C . Then all the infinitely many disjoint paths that start on this tail and end in X have to pass through the finite set S , a contradiction.

□

3. Statement of results, and open problems

Recall that $G = (V, E, \Omega)$ is an arbitrary fixed countable graph. Given two sets $A, B \subseteq V \cup \Omega$, we say that G *satisfies the Erdős-Menger conjecture for A and B* if $|G|$ contains a set \mathcal{P} of disjoint A - B paths and an A - B separator on \mathcal{P} .

Theorem 3.1. (Aharoni 1987)

G *satisfies the Erdős-Menger conjecture for all sets $A, B \subseteq V$.*

When proving Theorem 3.1 one may clearly assume that $Z := A \cap B = \emptyset$: having solved the problem for $G - Z$, just add every $z \in Z$ to X and, as a singleton path $\{z\}$, to \mathcal{P} , to obtain a solution for G . Interestingly, when A, B and X are allowed to contain ends, the assumption of $A \cap B = \emptyset$ is no longer just convenient for the proof, but the stronger assumption of $A \cap \overline{B} = \emptyset = \overline{A} \cap B$ is needed to make the conjecture true: a counterexample for $\overline{A} \cap B \neq \emptyset$ (Fig. 1) will be given at the end of this section. A more trivial counterexample, where G does not even contain an A - B path, would be a single ray R with $A = V(R)$ and B containing just the end of R (but this could be ‘repaired’ by relaxing the definition of an A - B path or of an A - B separator).

The following theorem, which is our main result, is therefore best possible in this sense:

Theorem 3.2. G *satisfies the Erdős-Menger conjecture for all sets $A, B \subseteq V \cup \Omega$ satisfying $A \cap \overline{B} = \emptyset = \overline{A} \cap B$.*

Note that A and B may contain uncountably many ends. We shall prove Theorem 3.2 in Sections 4 and 5, by first proving it for finite A directly and then using that to reduce the general case to Theorem 3.1.

The following consequence of Theorem 3.2 is a result of Stein [16]. Her proof, which is not easy, adapts Aharoni’s countable proof directly to ends, using the countability of the set $A \cup B$ implied by its discreteness:

Corollary 3.3. G *satisfies the Erdős-Menger conjecture for all sets $A, B \subseteq V \cup \Omega$ such that $A \cup B$ is a discrete subset of $|G|$.*

Proof. Making every point of $A \cap B$ into a trivial A - B path and putting it in the separator, we may assume that $A \cap B = \emptyset$. Then the discreteness of $A \cup B$ implies $A \cap \overline{B} = \emptyset = \overline{A} \cap B$, and the result follows from Theorem 3.2. \square

In our next corollary, A and B can again be uncountable:

Corollary 3.4. G *satisfies the Erdős-Menger conjecture for all sets $A, B \subseteq V \cup \Omega$ that are closed subsets of $|G|$.*

Proof. Make every point of $A \cap B$ into a trivial A - B path, put it in the separator, and apply Theorem 3.2 to $A' := A \setminus B$ and $B' := B \setminus A$. This can be done, since $\overline{A'} \subseteq A$ and $\overline{B'} \subseteq B$ by our assumption that A and B are closed, and hence $A' \cap \overline{B'} \subseteq A' \cap B = \emptyset = A \cap B' \supseteq \overline{A'} \cap B'$. \square

Let us finally sketch a counterexample to the Erdős-Menger conjecture for $\overline{A} \cap B \neq \emptyset$. Define a graph G as follows. Let $R = v_1 v'_1 v_2 v'_2 v_3 \dots$ be a ray. For all $i \in \mathbb{N}$, add new vertices a_i, a'_i, b_i and edges $a_i v_i, a'_i v'_i, v_i b_i, v'_i b_i$ (Fig. 1). Put $A := \{a_i \mid i \in \mathbb{N}\} \cup \{a'_i \mid i \in \mathbb{N}\}$ and $B := \{b_i \mid i \in \mathbb{N}\} \cup \{\beta\}$, where β is the unique end of G .




FIGURE 1. A counterexample for $\overline{A} \cap B \neq \emptyset$

Proposition 3.5. *The Erdős-Menger conjecture fails for these G, A, B .*

Proof. Suppose there is a set \mathcal{P} of disjoint A - B paths in G as well as an A - B separator X on \mathcal{P} . Let us show first that $b_i \in X$ for all $i \in \mathbb{N}$.

If $b_i \notin X$, then X meets both of the sets $\{a_i, v_i\}$ and $\{a'_i, v'_i\}$, say in $x \in P \in \mathcal{P}$ and $x' \in P' \in \mathcal{P}$, respectively. Not both P and P' can end in b_i ; we assume that $b_i \notin P'$. As clearly $v_i \in P$, the vertices of P' following x' (which include v_{i+1} and either v'_{i+1} or b_{i+1}) avoid X but separate a_{i+1} from B and A from b_{i+1} . Hence neither a_{i+1} nor b_{i+1} can lie on any path in \mathcal{P} other than P' , and the A - B path $a_{i+1} v_{i+1} b_{i+1}$ avoids X (a contradiction).

The fact that $b_i \in X$ for all i implies that X has no vertex outside B : any such vertex x would lie on a path $P \in \mathcal{P}$ ending in β , and then the final segment of P after x would be a tail of R meeting all the paths in \mathcal{P} that end in some b_i with i large enough (a contradiction). Hence $a_1 R$ is a path that starts in $a_1 \in A$ and ends in $\beta \in B$ but avoids X . \square

Proposition 3.5 shows that the condition $A \cap \overline{B} = \emptyset = \overline{A} \cap B$ in Theorem 3.2 cannot simply be dropped. The example of Figure 1 once led me to believe that the Erdős-Menger conjecture with ends might hold unconditionally for trees if $A \cup B \subseteq \Omega$, but the following example due to Stein (personal communication) refutes this. Let G be obtained from the binary tree T by grafting a new ray on to every vertex, let A contain the ends of these (countably many) new rays, and let B contain the (uncountably many) original ends of T . Then the Erdős-Menger conjecture fails for these G, A, B (and $\overline{A} \cap B \neq \emptyset$, in fact $\overline{A} \supseteq B$).

Although the original Erdős-Menger conjecture remains unproved for arbitrary uncountable graphs, one may ask whether in those cases where it is known to hold (see the Introduction) it generalizes to ends just as in the countable case. When A is countable (but allowed to contain ends as well as vertices) this can indeed be done. More generally, Theorem 3.2 is extended in [6] to arbitrary graphs in which A and B are separated by some countable set of vertices and ends.

Like the finite Menger theorem, the vertex version of the Erdős-Menger conjecture can be expressed also in terms of linking two non-adjacent vertices a and b rather than two vertex sets A and B . The conjecture then asks for a set of independent (rather than disjoint) a - b paths and an a - b separator on these paths that does not include a or b . This vertex version follows from our set version (take $A := N(a)$ and $B := N(b)$) and implies it (add a new vertex a adjacent to all of A and a new vertex b adjacent to all of B). Asking for a set of double rays joining two given ends α and β together with an α - β separator on these rays, however, yields a different problem that does not reduce as above to the set version treated in this paper, because ends do not have neighbours. The cardinality version of this other Menger-type problem for ends is treated by Halin [12] and by Polat [15] (who also considers sets A, B but allows his A - B paths to have common ends in A or B).

Finally, one might consider the Erdős-Menger problem topologically in the sense of [8, 9], asking for topological paths (ie. continuous images of the unit interval, without loss of generality 1-1) in $|G|$ between two sets $A, B \subseteq V \cup \Omega$ and a set X of vertices or ends ‘on’ these paths such that every topological A - B path in $|G|$ meets X . The following observations on this are due to Kühn (personal communication).

The graph G of Figure 2 satisfies $A \cap \overline{B} = \emptyset = \overline{A} \cap B$ but is a counterexample to the ‘topological’ Erdős-Menger conjecture as stated above. Indeed, suppose there is a set \mathcal{P} of topological A - B paths in $|G|$ and a topological A - B separator X on \mathcal{P} . Clearly, every path in \mathcal{P} is a finite path of length 2. Then all the points in X must be vertices in the middle (those of degree 4), and so the top ray and the bottom ray together with the unique end of G (to which both these rays converge) form a topological path from A to B that avoids X .




FIGURE 2. A counterexample to the topological version of the conjecture for $\overline{A} \cap \overline{B} \neq \emptyset$

The obvious way to avoid this counterexample would be to strengthen the disjointness requirement on A and B to $\overline{A} \cap \overline{B} = \emptyset$, a condition obviously violated in the graph of Figure 2. However, this condition implies also that the separator X in any Erdős-Menger system (\mathcal{P}, X) as provided by Theorem 3.2 separates A and B even topologically. Indeed, since the points of X lie on disjoint A - B paths, any end in $\overline{X} \setminus X$ would lie in $\overline{A} \cap \overline{B} = \emptyset$. So $\overline{X} \setminus X = \emptyset$, ie. X is closed in $|G|$. But this implies that any topological A - B path Q in $|G| \setminus X$ can be modified into a proper A - B path in $|G| \setminus X$, ie. one that contains no ends of G other than possibly its endpoints in A and B . (The proof of this is an easy adaptation of the proof of [9, Lemma 3.4].) We thus have the following strengthening of Theorem 3.2 for sets A, B with disjoint closures (even when G is uncountable):

Theorem 3.6. *If $\overline{A} \cap \overline{B} = \emptyset$, then any solution in $|G|$ to the Erdős-Menger conjecture for $A, B \subseteq V \cup \Omega$ also solves the topological Erdős-Menger conjecture in $|G|$ for A and B . \square*

If the imposition of the condition of $A \cap \overline{B} = \emptyset = \overline{A} \cap B$ in Theorem 3.2 seemed unnecessarily strong at least for some natural classes of graphs, then imposing $\overline{A} \cap \overline{B} = \emptyset$ in Theorem 3.6 will seem worse: there are many examples of G, A, B that do not satisfy $\overline{A} \cap \overline{B} = \emptyset$ but where the topological conjecture can be proved nonetheless. So as before, the questions of how this condition can be weakened, for which classes of graphs it can be avoided altogether, and how the topological Erdős-Menger conjecture relates to the standard version when $\overline{A} \cap \overline{B} \neq \emptyset$ remain open.

4. Alternating paths between ends

In this section we adapt standard alternating path techniques, such as used by Gallai [11] in his constructive proof of the finite Menger theorem, to our infinite setting. The exposition will concentrate on the infinite aspects that are new, and skip some of the finite details; these can be found in the rendering of Gallai's proof in [7].

Let H be a subgraph of G , let $S, T \subseteq V(H) \cup \Omega$ be disjoint, and let \mathcal{P} be a set of disjoint S - T paths in \overline{H} . We write $|\mathcal{P}|$ for the union of all the paths in \mathcal{P} including their endpoints; thus,

$$|\mathcal{P}| := \bigcup \{ \overline{P} \mid P \in \mathcal{P} \}.$$

An *alternating walk* with respect to \mathcal{P} is a finite sequence $\mathcal{W} = (W_1, \dots, W_n)$, with n odd, of paths in \overline{H} satisfying the following conditions:

- (A1) the first but no other point of W_1 lies in $S \setminus |\mathcal{P}|$, the last point of W_n lies in $|\mathcal{P}|$ or in $T \setminus |\mathcal{P}|$, and no point of W_n other than the last lies in T ;

- (A2) except for the endpoints mentioned in (i), all W_i with i odd have both endpoints in $|\mathcal{P}|$ and are otherwise contained in $H \setminus |\mathcal{P}|$;
- (A3) every W_i with i even is a non-trivial segment of some \overline{P} with $P \in \mathcal{P}$ in reverse order;
- (A4) distinct W_i are disjoint except possibly in their endpoints, and for all $i = 1, \dots, n-1$ the last point of W_i is the first point of W_{i+1} .

Note that there are only finitely many W_i , but that the W_i themselves may be infinite: they can be rays, inverse rays, or even double rays with ends in $S \cup T$.

The concatenation W of the paths W_1, \dots, W_n above is a topological path in \overline{H} , ie. a continuous (though not necessarily 1-1) image in \overline{H} of the unit interval $[0, 1]$. We do not always distinguish \mathcal{W} from W notationally, and may thus speak of the ‘first point’ of \mathcal{W} in a given (closed) set $D \subseteq \overline{H}$ etc.

Another set \mathcal{Q} of disjoint S - T paths in \overline{H} will be said to *exceed* \mathcal{P} if $|\mathcal{P} \cap S \subsetneq |\mathcal{Q} \cap S$ and $|\mathcal{P} \cap T \subsetneq |\mathcal{Q} \cap T$.

Lemma 4.1. *If \overline{H} contains an alternating walk with respect to \mathcal{P} that ends in $T \setminus |\mathcal{P}|$, then \overline{H} contains a set of disjoint S - T paths exceeding \mathcal{P} .*

Proof. Let $\mathcal{W} = (W_1, \dots, W_n)$ be an alternating walk as stated. Let H' be the subgraph of H whose edge set is the symmetric difference between the edge sets of $\bigcup \mathcal{W}$ and $\bigcup \mathcal{P}$, and whose vertices are those incident with these edges.

Local inspection of how taking this symmetric difference affects the degree of a given vertex on \mathcal{P} or \mathcal{W} shows that every vertex of H' outside $S \cup T$ has degree 2, while those in $S \cup T$ have degree 1. So the components of H' are either paths, or finite cycles avoiding $S \cup T$. Moreover since n is finite, every ray or inverse ray contained in some $P \in \mathcal{P}$ or in some W_i has a tail that either lies in H' or avoids H' . Together with (A4), the above implies that every vertex or end in S that is a first point of some $P \in \mathcal{P}$ or of W_1 , as well as every vertex or end in T that is a last point of some $P \in \mathcal{P}$ or of W_n , is also an endpoint of a path that is a component of H' .

Finally, any path Q that is a component of H' traverses all its edges in their original direction as induced by some $P \in \mathcal{P}$ or some W_i with i odd. Therefore Q has its first but no other point in S and its last but no other point in T , and hence is an S - T path.

The set of components of H' therefore includes a set \mathcal{Q} of disjoint S - T paths in \overline{H} whose first points include $S \cap |\mathcal{P}|$ as well as the first point of W_1 , and whose last points include $T \cap |\mathcal{P}|$ as well as the last point of W_n . Thus, \mathcal{Q} exceeds \mathcal{P} . \square

Lemma 4.2. *If $S \cap \overline{T} = \emptyset = \overline{S} \cap T$, then in \overline{H} there is either an S - T separator on \mathcal{P} or a set of disjoint S - T paths exceeding \mathcal{P} .*

Proof. Unless otherwise stated, all alternating walks in this proof will be with respect to \mathcal{P} . Let us define a potential S - T separator X on \mathcal{P} by selecting one point x_P from every path $P \in \mathcal{P}$, as follows. If P has a last vertex that lies on some alternating walk, let x_P be that vertex. If P has a vertex on an alternating walk but no last such vertex, we let x_P be the last point of P (which will be an end in T). If P has no vertex on an alternating walk, let x_P be the first point of P . We show that if

$$X := \{x_P \mid P \in \mathcal{P}\}$$

is not an S - T separator in \overline{H} , then \overline{H} contains an alternating walk with respect to \mathcal{P} that ends in $T \setminus |\mathcal{P}|$. The assertion then follows from Lemma 4.1.

Suppose there is a path Q in \overline{H} that starts in S , ends in T , and satisfies $\overline{Q} \cap X = \emptyset$. Since $S \cap \overline{T} = \emptyset = \overline{S} \cap T$, clearly Q has a last point in S and after this a first point in T , so we may assume that Q is an S - T path. If $\overline{Q} \cap |\mathcal{P}| = \emptyset$, then $\mathcal{W} = (Q)$ is the desired alternating walk (and $\mathcal{P} \cup \{Q\}$ is a set of disjoint S - T paths exceeding \mathcal{P}). So we assume that $\overline{Q} \cap |\mathcal{P}| \neq \emptyset$. Let us show that Q has a first point in $|\mathcal{P}|$. If not, then the first point of Q is an end $\sigma \in S \setminus |\mathcal{P}|$, and Q contains an (inverse) ray $R \in \sigma$ that has infinitely many vertices in $|\mathcal{P}|$. But only finitely many of these vertices lie on any one $P \in \mathcal{P}$, since otherwise P would have a tail equivalent to $R \in \sigma \notin |\mathcal{P}|$, a contradiction. So R meets infinitely many distinct paths $P \in \mathcal{P}$, but each of them in only finitely many vertices. Since each of these P has an endpoint in T , we thus have infinitely many disjoint R - T paths in \overline{H} and hence $R \in \sigma \in \overline{T}$ by Lemma 2.1. As $\sigma \in S$, this contradicts our assumption that $S \cap \overline{T} = \emptyset$. So Q does indeed have a first point in $|\mathcal{P}|$; let us call this point p .

Unless p is the first point of Q , the initial segment Qp of Q ending in p is an alternating walk. Moreover, if p is an end in T , then $p = x_P$ for the path $P \in \mathcal{P}$ containing p , because QpP is an alternating walk containing arbitrarily late vertices of P . Hence in either case Q also meets the set $|\mathcal{P}'|$, where

$$\mathcal{P}' := \{Px_P \mid P \in \mathcal{P}\}$$

is the set of initial segments of the paths $P \in \mathcal{P}$ ending in x_P . Let us show that Q has a last point in $|\mathcal{P}'|$. If not, then Q has a (forward) tail R with infinitely many vertices in $|\mathcal{P}'|$; let $\tau \in T$ be the end containing R . Suppose first that R shares infinitely many vertices from $|\mathcal{P}'|$ with some fixed path $P \in \mathcal{P}$. Then P has a tail with infinitely many vertices in $V(R) \cap |\mathcal{P}'|$, which must be a forward tail of P because $R \in \tau \in T$. So P has arbitrarily late vertices in $V(R) \cap |\mathcal{P}'|$, implying $x_P = \tau$. But then Q has its last point τ in X , contrary to our assumption that $\overline{Q} \cap X = \emptyset$. Hence of the infinite set $V(R) \cap |\mathcal{P}'|$ only finitely many points lie on any one $P \in \mathcal{P}$, so R meets infinitely many different $P \in \mathcal{P}$. But now \overline{H} contains infinitely many disjoint paths from S to R , implying $R \in \tau \in \overline{S}$ by Lemma 2.1. As $\tau \in T$, this contradicts our assumption that $\overline{S} \cap T = \emptyset$. So Q does indeed have a last point in $|\mathcal{P}'|$; let us call this point q .

Let P be the path in \mathcal{P} containing q , and let $x := x_P$. As Q avoids X we have $q \neq x$, so q precedes x on P (as $q \in |\mathcal{P}'|$) and is not the last point of Q . In particular, $x = x_P$ is not the first point of P , and so there is an alternating walk \mathcal{W} ending on $\dot{q}P$ (ie. on P after q , for example at x). Let y be the first point on $\dot{q}P$ of the topological path W corresponding to \mathcal{W} ; then $W' := WyPq$ defines an alternating walk from $S \setminus |\mathcal{P}|$ to q . If W' meets qQ only in q , let $W'' := W' \cup qQ$; if W' meets $\dot{q}Q$, say first in z , let $W'' := W'zQ$. In either case W'' defines an alternating walk with respect to \mathcal{P}' .

But no alternating walk with respect to \mathcal{P}' can meet $|\mathcal{P}|$ outside $|\mathcal{P}'|$: if it did, it would do so in some first point r , and then its initial segment ending in r would be an alternating walk with respect to \mathcal{P} , putting r in $|\mathcal{P}'|$ by definition of $|\mathcal{P}'|$. Therefore W'' is an alternating walk also with respect to \mathcal{P} , and the last point of W'' (which is also the last point of $\dot{q}Q \neq \emptyset$) cannot lie in $|\mathcal{P}|$. Hence W'' ends in $T \setminus |\mathcal{P}|$, as required. \square

Corollary 4.3. *If S is a finite set of vertices, then \overline{H} contains a set \mathcal{P} of disjoint S - T paths and an S - T separator on \mathcal{P} .*

Proof. Making every point in $S \cap T$ into a trivial S - T path and putting it in the separator, we may assume that $S \cap T = \emptyset$. Since S is finite and consists of vertices only, this implies $S \cap \overline{T} = \emptyset = \overline{S} \cap T$. Now let \mathcal{P} be a set of disjoint S - T paths in \overline{H} , as many as possible. This set cannot be exceeded, so the assertion follows from Lemma 4.2. \square

5. The end-to-vertex reduction

In this section we shall use Corollary 4.3 to reduce the end version of the Erdős-Menger conjecture to its vertex version. Since the vertex version is currently known only for countable graphs, our proof will make free use of the countability of G . This will simplify matters considerably but could, I believe, be avoided.

We begin with a lemma from Stein [16]. Let T be a finite set of vertices in a graph H . A T -path, for the purpose of this paper, is any path whose endvertices lie in T , whose inner vertices lie outside T , and which has at least one inner vertex. Paths P_1, \dots, P_k are said to be *disjoint outside* some given $Q \subseteq H$ if $P_i \cap P_j \subseteq Q$ whenever $i \neq j$.

Lemma 5.1. *Let H be a graph, let $T \subseteq V(H)$ be finite, and let $k \in \mathbb{N}$. Then H has a finite subgraph H' containing T such that for every T -path $Q = s \dots t$ in H that meets $H - H'$ there are k distinct T -paths from s to t in H' that are disjoint outside Q (Fig. 3).*

Proof. We apply induction on k . For $k = 0$ the lemma holds with $H' := H[T]$. For the induction step, let $H' = H'_k$ satisfy the lemma for k . For every two vertices $u, v \in H'_k$ such that H contains a $V(H'_k)$ -path from u to v , add such a path to H'_k and call the resulting graph H'_{k+1} .




FIGURE 3. T -paths P_1, \dots, P_k that are disjoint outside Q

To show that $H' = H'_{k+1}$ satisfies the lemma for $k+1$, let $Q = s \dots t$ be given as stated. Since Q meets $H - H'_{k+1} \subseteq H - H'_k$, the induction hypothesis provides us with k distinct T -paths P_1, \dots, P_k from s to t in H'_k that are disjoint outside Q . Moreover, by the definition of H'_{k+1} , for every $V(H'_k)$ -path $u \dots v$ contained in Q there exists a $V(H'_k)$ -path from u to v in H'_{k+1} . The union of these latter paths with $Q \cap H'_k$ is a connected subgraph of H'_{k+1} containing s and t , which meets H'_k (and hence P_1, \dots, P_k) only on Q ; let P_{k+1} be an s - t path in this connected graph. Then every edge of P_{k+1} with both endvertices in H'_k is an edge of Q . Since $Q \cap H'_k$ does not contain an s - t path, this means that P_{k+1} has a vertex outside H'_k and therefore differs from P_1, \dots, P_k . \square

Our next lemma is the essential step in our reduction of the end version of the Erdős-Menger conjecture to its vertex version: it replaces the set $A \subseteq V \cup \Omega$ with a set A' consisting only of vertices, and can then be repeated to do the same for B .

Lemma 5.2. *Let $A, B \subseteq V \cup \Omega$ be such that $A \cap \overline{B} = \emptyset = \overline{A} \cap B$. Then G has a minor $G' = (V', E', \Omega')$ with subsets $A' \subseteq V'$ and $B' \subseteq V' \cup \Omega'$ such that the following two assertions hold:*

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

Proof. Choose an enumeration $V = \{v_1, v_2, \dots\}$ of the vertices of G , and for every $i \in \mathbb{N}$ put $S_i := \{v_1, \dots, v_i\}$. For $i = 0, 1, \dots$ in turn, define \mathcal{C}_i as the set of all the components C of $G - S_i$ that satisfy the following requirements:

- (C1) $\overline{C} \cap B = \emptyset$;
- (C2) $\overline{C} \cap A \cap \Omega \neq \emptyset$;
- (C3) $C \cap D = \emptyset$ for every $D \in \mathcal{C}_j$ with $j < i$.

Finally, put $\mathcal{C} := \bigcup_{i \in \mathbb{N}} \mathcal{C}_i$.

By (C3), the subgraphs $C \in \mathcal{C}$ are pairwise disjoint. In fact, they have disjoint closures \overline{C} : since every $N(C) \subseteq S_i$ is finite, no end of G can have rays in two distinct $C \in \mathcal{C}$. When we delete and contract parts of G to form G' , we shall do this independently in the various $C \in \mathcal{C}$.

Although formally S_i will often have vertices s in components $D \in \mathcal{C}_j$ with $j < i$, no such s can have a neighbour in any $C \in \mathcal{C}_i$, because the neighbours of $s \in D$ lie either in D (which is disjoint from C) or in $S_j \subseteq S_i$. Since our purpose in considering a component $C \in \mathcal{C}_i$ will lie in finding S_i - A paths in the graph $G[C \cup S_i]$, such vertices s will thus be useless. To simplify matters, we shall therefore use

$$S_C := N(C) \subseteq S_i$$

instead of S_i to separate off C , and consider the graphs

$$G_C := G[C \cup S_C]$$

instead of $G[C \cup S_i]$. Then

$$G_C \cap D = \emptyset \text{ for all distinct } C, D \in \mathcal{C}. \quad (1)$$

Indeed, we already observed that $C \cap D = \emptyset$. Moreover, we just showed that S_C cannot meet any $D \in \mathcal{C}_j$ with $j < i$. And S_C cannot meet any $D \in \mathcal{C}_j$ with $j \geq i$, because $S_C \subseteq S_i \subseteq S_j$. This completes the proof of (1).

Although every $C \in \mathcal{C}$ avoids B , let us remember that the sets S_C may meet B . For each $C \in \mathcal{C}$, put

$$A_C := A \cap \overline{G_C}.$$

Let us show the following:

$$\text{Every end } \alpha \in A \cap \Omega \text{ lies in } A_C \text{ for some } C \in \mathcal{C}. \quad (2)$$

As $A \cap \overline{B} = \emptyset$ by assumption, α has a neighbourhood $\widehat{C}(S, \alpha)$ in $|G|$ that avoids B . If i is large enough that $S \subseteq S_i$, the component C' of $G - S_i$ to which α belongs is contained in C and hence satisfies (C1) and (C2). Therefore $C' \in \mathcal{C}_i$ unless C' fails to satisfy (C3), in which case C' meets some $D \in \mathcal{C}_j$ with $j < i$. But then $C' \subseteq D$ because $S_j \subseteq S_i$, giving $\alpha \in A_D$. This completes the proof of (2).

We now describe, independently for distinct $C \in \mathcal{C}$, which parts of C are to be deleted or contracted to form G' . So fix $C \in \mathcal{C}$. We first use Corollary 4.3 to find in $\overline{G_C}$ a set \mathcal{P}_C of disjoint S_C - A_C paths together with an S_C - A_C separator X_C on \mathcal{P}_C (Fig. 4). Let us write X_C as $X_C = U_C \cup O_C$, where

By (3) and the definition of H_C , we have $\overline{H_C} \cap A \subseteq U_C \cup O_C = X_C$. Since O_C is finite, we can extend $U_C \cup S_C$ to a finite set $T_C \subseteq V(H_C)$ that separates the ends in O_C pairwise in G . For each $\alpha \in O_C$ let $C_\alpha \subseteq H_C \cap C$ be the component of $G - T_C$ to which α belongs, and put $H_\alpha := G[C_\alpha \cup T_C]$. Let H'_α be the finite subgraph of H_α containing T_C which Lemma 5.1 provides for $k := |T_C| + 2$, and let $D_\alpha \subseteq C_\alpha$ be the component of $G - H'_\alpha$ to which α belongs. Write $\mathcal{D}_2(C) := \{D_\alpha \mid \alpha \in O_C\}$ and $\mathcal{D}_2 := \bigcup_{C \in \mathcal{C}} \mathcal{D}_2(C)$. Finally put $\mathcal{D} := \mathcal{D}_1 \cup \mathcal{D}_2$, noting that $D \cap D_\alpha = \emptyset$ for all $D \in \mathcal{D}_1$ and $D_\alpha \in \mathcal{D}_2$, by (4).

In order to define G' , put

$$U := \bigcup_{C \in \mathcal{C}} U_C \quad \text{and} \quad V_{AB} := U \cap B.$$

Let \tilde{G} be obtained from G by deleting $\bigcup \mathcal{D}_1$ and contracting every $D_\alpha \in \mathcal{D}_2$ to a single vertex a_α . Let

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

and put

$$\tilde{A} := \left(A \setminus \bigcup_{C \in \mathcal{C}} A_C \right) \cup U \cup A^* \stackrel{(1,2)}{\subseteq} V(\tilde{G}).$$

Finally, let

$$G' := \tilde{G} - V_{AB} \quad \text{and} \quad A' := \tilde{A} \setminus V_{AB}.$$

Then for $Z := V_{AB} \cup \bigcup_{D \in \mathcal{D}} V(D)$ we have

$$G - Z = G \cap G' = G' - A^*.$$

An important property of G' is that the ends of G in B correspond closely to ends of G' . To establish this correspondence formally, we begin with the following observation:

$$\text{Every ray } R \text{ in an end } \beta \in B \text{ has a tail in } G - Z. \quad (6)$$

Indeed, as for every $D \in \mathcal{D}$ the set $N(D)$ of its neighbours in G is finite, each D can meet R in at most finitely many vertices; recall that $\overline{D} \subseteq \overline{C}$ for some $C \in \mathcal{C}$, and hence $\beta \notin \overline{D}$ by (C1). It remains to show that R cannot meet infinitely many $D \in \mathcal{D}$ and has only finitely many vertices in V_{AB} . Recall that distinct $D \in \mathcal{D}$ are disjoint, and that $\overline{D} \cap A \neq \emptyset$ by definition of D . Hence if R meets infinitely many $D \in \mathcal{D}$, we can find infinitely many disjoint R - A paths in G . Similarly if R has infinitely many vertices in $V_{AB} \subseteq U$, we can find infinitely many disjoint R - A paths in G by (5) and (1). But then $\beta \in \overline{A}$ by Lemma 2.1, contrary to our assumption that $\overline{A} \cap B = \emptyset$. This completes the proof of (6).

Let R_1, R_2 be two rays in $G \cap G'$, and assume that the end of R_1 in G lies in B . Then R_1 and R_2 are equivalent in G if and only if they are equivalent in G' . (7)

To prove (7), suppose first that R_1 and R_2 are equivalent in G , ie. belong to the same end $\beta \in B$. Then G contains infinitely many disjoint R_1 – R_2 paths P_1, P_2, \dots . Since $\beta \notin \overline{A}$, only finitely many $D \in \mathcal{D}$ can have a vertex in $\bigcup_i P_i$ (as in the proof of (6)), and each of these D meets only finitely many P_i , because $N(D)$ is finite and $R_1, R_2 \subseteq G - D$. Similarly, $\beta \notin \overline{A}$ implies that only finitely many P_i meet $V_{AB} \subseteq U$. Hence all but finitely many of the P_i lie in $G \cap G'$, showing that R_1 and R_2 are equivalent also in G' .

Conversely, if R_1 and R_2 are joined in G' by infinitely many disjoint paths, we can replace any vertices $a_\alpha \in V' \setminus V = A^*$ on these paths by finite paths in D_α to obtain infinitely many disjoint R_1 – R_2 paths in G . (Alternatively, one can observe as in (6) that only finitely many of those R_1 – R_2 paths in G' can have vertices in A^* .) This completes the proof of (7).

We can now define our correspondence between the ends in B and certain ends of G' , and then define B' . For every end $\beta \in B$ there is by (6) an end $\beta' \in \Omega' = \Omega(G')$ such that $\beta \cap \beta' \neq \emptyset$. By (7), this end β' is unique and the map $\beta \mapsto \beta'$ is injective. Let

$$B' := \{ \beta' \mid \beta \in B \cap \Omega \} \cup (B \cap V) \setminus V_{AB} \underset{(C1)}{\subseteq} \Omega' \cup V' \setminus A'.$$

Let us prove assertion (i) of the lemma (with closures taken in $|G'|$). We trivially have $A' \cap \overline{B'} = \emptyset$, because $A' \subseteq V'$ and $A' \cap B' = \emptyset$. To show that $\overline{A'} \cap B' = \emptyset$, consider an end $\beta' \in B'$. The corresponding end $\beta \in B$ has a neighbourhood $\widehat{C}(S, \beta)$ in $|G|$ that avoids A . By (5), this C has only finitely many vertices in U and meets only finitely many $D \in \mathcal{D}$ (as in the proof of (6)). Adding to $S \setminus Z$ the finite set

$$V' \cap V(C) \cap (U \cup \bigcup \{ N(D) \mid D \in \mathcal{D}; D \cap C \neq \emptyset \})$$

then yields a finite set $S' \subseteq V'$ such that the neighbourhood $\widehat{C}'(S', \beta')$ of β' in $|G'|$ does not meet A' , as desired.

We now prove assertion (ii) of the lemma. Suppose that G' contains a set \mathcal{P}' of disjoint A' – B' paths and an A' – B' separator X' on \mathcal{P}' . Put

$$\tilde{\mathcal{P}} := \mathcal{P}' \cup \{ (x) \mid x \in V_{AB} \},$$

where (x) denotes the trivial \tilde{A} – B path with vertex x .

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 $a \in \tilde{A}$ of \tilde{P} lies in A we leave \tilde{P} unchanged, ie. set $P := \tilde{P}$. If $a \in U \setminus A$, pick $C \in \mathcal{C}$ with $a \in U_C$, and let P be the union of \tilde{P} with an A_C – H_C path in G_C that ends in a ; this can be done

disjointly for different $\tilde{P} \in \tilde{\mathcal{P}}$ if we use the paths from (5) and remember that distinct $C \in \mathcal{C}$ have disjoint closures. Moreover, the A_C-H_C path concatenated with \tilde{P} in this way has all its vertices other than a outside G' and V_{AB} , so it contains no other vertices of $\tilde{\mathcal{P}}$. 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 any new initial segment of P lies in a subgraph $D \in \mathcal{D}$ of G , which avoids B because the $C \in \mathcal{C}$ containing it satisfies (C1).

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 $\beta' \in B' \cap \Omega'$ then its tail $\tilde{P} - a \subseteq P \subseteq G$ is equivalent in G' to some ray in $\beta' \cap \beta$, by definition of β' . By (7) 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 , since the corresponding paths \tilde{P} have different endpoints β' by assumption.

We still need an A - B separator X on \mathcal{P} . Let X be obtained from $X' \cup V_{AB}$ by replacing every end $\beta' \in X'$ with the corresponding end $\beta \in B$, and replacing every $a_\alpha \in X' \cap A^*$ 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} .

To show that X separates A from B in G , suppose there exists a path $Q \subseteq G - X$ that starts in A and ends in B . Note that Q avoids V_{AB} , since $V_{AB} \subseteq X$. Our aim is to turn Q into an A' - B' path $Q' \subseteq G' - X'$, with a contradiction. If Q meets $\bigcup \mathcal{D}_1$ it has a last vertex there by (6), and its next vertex a lies in $U \setminus V_{AB} \subseteq A'$; we then define Q' (for the time being) 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 $a \in A \cap A'$ (in which case we put $Q' := Q$), or Q starts in an end $\alpha \in O_C$ for some $C \in \mathcal{C}$; we then make $a := a_\alpha$ the starting vertex of Q' and continue Q' along Q , beginning with the last D_α - G' edge on Q . Then our assumption of $\alpha \notin X$ (by the choice of Q) implies that $a \notin X'$, by the definition of X . Thus in all three cases, Q' is now a path that avoids X' , and which starts at a vertex $a \in A'$ and from there continues with a tail of Q that avoids $\bigcup \mathcal{D}_1$.

However, Q' may still not be a path in G' , since it can meet some components $D_\alpha \in \mathcal{D}_2$ with $a_\alpha \neq a$. Indeed, Q might use such segments in D_α to bypass X . And we may not be able to turn Q' into our desired A' - B' path in $G' - X'$ simply by replacing those Q -segments in D_α with the vertex $a_\alpha \in V'$, because it may happen that $a_\alpha \in X'$. Using Lemma 5.1, however, we shall be able to replace any segments of Q' that meet some $D_\alpha \in \mathcal{D}_2$ with paths through the corresponding graph $H'_\alpha \subseteq G \cap G'$ that avoid X' . Since Q (and hence Q') has by (6) a tail that avoids all $D \in \mathcal{D}_2$, these modifications will only affect a finite initial segment of Q' . So all these modifications will turn Q' into a (walk

that can be pruned to a) path that starts at $a \in A'$, and then runs through $G \cap G'$ until it ends at the last vertex b of Q or in the forward end β of Q (Fig. 5). If we view Q' as a path in G' , its last point is either b or β' in B' , yielding the desired contradiction.




FIGURE 5. Modifying $Q' - a$ into a path in $G \cap G'$

So consider a segment sQt of Q' that meets some $D_\alpha \in \mathcal{D}_2$ with $a_\alpha \neq a$. By definition of D_α , we may choose s and t so that sQt is a T_C -path in H_α , where C is the unique element of \mathcal{C} containing D_α . By definition of the graph $H'_\alpha \subseteq H_\alpha$ (which is a subgraph of \tilde{G} by (1), ie. no parts of H'_α were deleted or contracted during the treatment of any other $C \in \mathcal{C}$), there are $|T_C| + 2$ paths from s to t in H'_α that are disjoint outside sQt . But H'_α contains at most $|T_C| + 1$ vertices from $X' \cup V_{AB}$: since these lie on disjoint paths \tilde{P} starting in \tilde{A} , and the only point in \tilde{A} not separated in \tilde{G} from H'_α by T_C is a_α (by definition of T_C), all but at most one of these paths \tilde{P} meet T_C on their way from \tilde{A} to H'_α . So one of our $|T_C| + 2$ distinct $s-t$ paths in H'_α that are disjoint outside sQt avoids $X' \cup V_{AB}$, and we can use this path to replace sQt on Q' . \square

Proof of Theorem 3.2. Apply Lemma 5.2 twice, to replace first A and then B (or B') by a set of vertices. Then apply Aharoni's Theorem 3.1. \square

Acknowledgement

Thanks are due to Henning Bruhn for pointing out a subtle technical error in an earlier draft of the proof of Lemma 5.2.

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* **31** (1987), 303–313.
3. R. Aharoni, A few remarks on a conjecture of Erdős on the infinite version of Menger's theorem, in (R.L. Graham & J. Nešetřil, eds.): *The Mathematics of Paul Erdős II*, Springer-Verlag Heidelberg 1997.
4. R. Aharoni & R. Diestel, Menger's theorem for a countable source set, *Combinatorics, Probability and Computing* **3** (1994), 145–156.
5. R. Aharoni, C.St.J.A. Nash-Williams & S. Shelah, A general criterion for the existence of transversals, *Proc. London Math. Soc.* **47** (1983), 43–68.
6. H. Bruhn, R. Diestel & M. Stein, The Erdős-Menger conjecture with ends in arbitrary graphs with countable separators, submitted.
7. R. Diestel, *Graph Theory* (2nd edition), Springer-Verlag 2000.
<http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html>
8. R. Diestel & D. Kühn, On infinite cycles I, *Combinatorica* (to appear).
9. R. Diestel & D. Kühn, On infinite cycles II, *Combinatorica* (to appear).
10. R. Diestel & D. Kühn, Graph-theoretical versus topological ends of graphs, *J. Combin. Theory B* (to appear).
11. T. Grünwald, Ein neuer Beweis eines Mengerschen Satzes, *J. London Math. Soc.* **13** (1938), 188–192.
12. R. Halin, A note on Menger's theorem for infinite locally finite graphs, *Abh. Math. Sem. Univ. Hamburg* **40** (1974), 111–114.
13. H.A. Jung, Connectivity in infinite graphs, in (L. Mirsky, ed.) *Studies in Pure Mathematics*, Academic Press 1971.
14. C.St.J.A. Nash-Williams, Infinite graphs, a survey, *J. Combin. Theory B* **3** (1967), 40–46.
15. N. Polat, A Mengerian theorem for infinite graphs with ideal points, *J. Combin. Theory B* **51** (1991), 248–255.
16. M. Stein, *Die Erdős-Menger'sche Vermutung für abzählbare Graphen mit Enden*, Diplomarbeit Universität Hamburg 2002.

Reinhard Diestel
Mathematisches Seminar der
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany

diestel@math.uni-hamburg.de

Version 10 January 2003