MENGER'S THEOREM FOR INFINITE GRAPHS

RON AHARONI AND ELI BERGER

Abstract. We prove that Menger's theorem is valid for infinite graphs, in the following strong version: let A and B be two sets of vertices in a possibly infinite digraph. Then there exist a set $\mathcal{P}$ of disjoint A–B paths, and a set S of vertices separating A from B, such that S consists of a choice of precisely one vertex from each path in $\mathcal{P}$. This settles an old conjecture of Erdős.

1. History of the problem

In 1931 Dénes König [17] proved a min-max duality theorem on bipartite graphs:

Theorem 1.1. *In any finite bipartite graph, the maximal size of a matching equals the minimal size of a cover of the edges by vertices.*

Here a *matching* in a graph is a set of disjoint edges, and a *cover* (of the edges by vertices) is a set of vertices meeting all edges. This theorem was the culmination of a long development, starting with a paper of Frobenius in 1912. For details on the intriguing history of this theorem, see [19]. Four years after the publication of König's paper Phillip Hall [16] proved a result which he named "the marriage theorem". To formulate it, we need the following notation: given a set A of vertices in a graph, we denote by $N(A)$ the set of its neighbors.

Theorem 1.2. *In a finite bipartite graph with sides M and W there exists a marriage of M (that is, a matching meeting all vertices of M) if and only if $|N(A)| \geq |A|$ for every subset A of M.*

The two theorems are closely related, in the sense that they are easily derivable from each other. In fact, König's theorem is somewhat stronger, in that the derivation of Hall's theorem from it is more straightforward than vice versa.

At the time of publication of König's theorem, a theorem generalizing it considerably was already known.

Definition 1.3. Let X, Y be two sets of vertices in a digraph D. A set S of vertices is called *X–Y-separating* if every X–Y-path meets S, namely if the deletion of S severs all X–Y-paths.

Note that, in particular, S must contain $X \cap Y$.

The research of the first author was supported by grant no. 780-04 of the Israel Science Foundation, by GIF grant no. 2006311, by the Technion's research promotion fund, and by the Discont Bank chair.

The research of the second author was supported by the National Science Foundation, under agreement No. DMS-0111298. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the view of the National Science Foundation.

Notation 1.4. The minimal size of an X–Y-separating set is denoted by $\sigma(X,Y)$. The maximal size of a family of vertex-disjoint paths from X to Y is denoted by $\nu(X,Y)$.

In 1927 Karl Menger [21] published the following:

Theorem 1.5. *For any two sets A and B in a finite digraph there holds:*

$$\sigma(A,B) = \nu(A,B) \ .$$

This was probably the first casting of a combinatorial result in min-max form. There was a gap in Menger's proof: he assumed, without proof, the bipartite case of the theorem, which is Theorem 1.1. This gap was filled by König. Since then other ways of deriving Menger's theorem from König's theorem have been found, see, e.g., [1].

Soon thereafter Erdős, who was König's student, proved that, with the very same formulation, the theorem is also valid for infinite graphs. This appeared in König's book [18], the first book published on graph theory. The idea of the proof is this: take a maximal family $\mathcal{P}$ of A–B-disjoint paths. The set $S = \bigcup\{V(P) : \ P \in \mathcal{P}\}$ is then A–B-separating, since an A–B-path avoiding it could be added to $\mathcal{P}$, contradicting the maximality of $\mathcal{P}$. Since every path in $\mathcal{P}$ is finite, if $\mathcal{P}$ is infinite then $|\mathcal{P}| = |S|$. Since $\nu(A,B) \geq |\mathcal{P}|$ and $\sigma(A,B) \leq |S|$, this implies the non-trivial inequality $\nu(A,B) \geq \sigma(A,B)$ of the theorem. If $\mathcal{P}$ is finite, then so is S. The size of families of disjoint $A - B$ paths is thus finitely bounded (in fact, bounded by $|S|$), and hence there exists a finite family of maximal cardinality of disjoint A–B paths. In this case one can apply one of many proofs known for the finite case of the theorem (see, e.g., Theorem 4.7 below, or [14]).

Of course, there is some "cheating" here. The separating set produced in the case that $\mathcal{P}$ is infinite is obviously too "large". In the finite case the fact that $|S| = |\mathcal{P}|$ implies that there is just one S-vertex on each path of $\mathcal{P}$, while in the infinite case the equality of cardinalities does not imply this. Erdős conjectured that, in fact, the same relationship between S and $\mathcal{P}$ can be obtained also in the infinite case. Since it is now proved, we state it as a theorem:

Theorem 1.6. *Given two sets of vertices, A and B, in a (possibly infinite) digraph, there exists a family $\mathcal{P}$ of disjoint A–B-paths, and a separating set consisting of the choice of precisely one vertex from each path in $\mathcal{P}$.*

The earliest reference in writing to this conjecture is [29] (Problem 8, p. 159. See also [22]).

The first to be tackled was of course the bipartite case, and the first breakthrough was made by Podewski and Steffens [27], who proved the countable bipartite case of the conjecture, namely the countable case of König's theorem. That paper established some of the basic concepts that were used in later work on the conjecture, and also set the basic approach: introducing an a-symmetry into the problem. In the conjecture (now theorem) the roles of A and B are symmetrical; the proof in [27] starts with asking the question of when can a given side of a bipartite graph be matched into the other side, namely the problem of extending Hall's theorem to the infinite case. Known as the "marriage problem", this question was open since the publication of Hall's paper, and Podewski and Steffens solved its countable case. Around the same time, Nash-Williams formulated two other necessary criteria for matchability (the existence of marriage), and he [24, 25] and Damerell and Milner

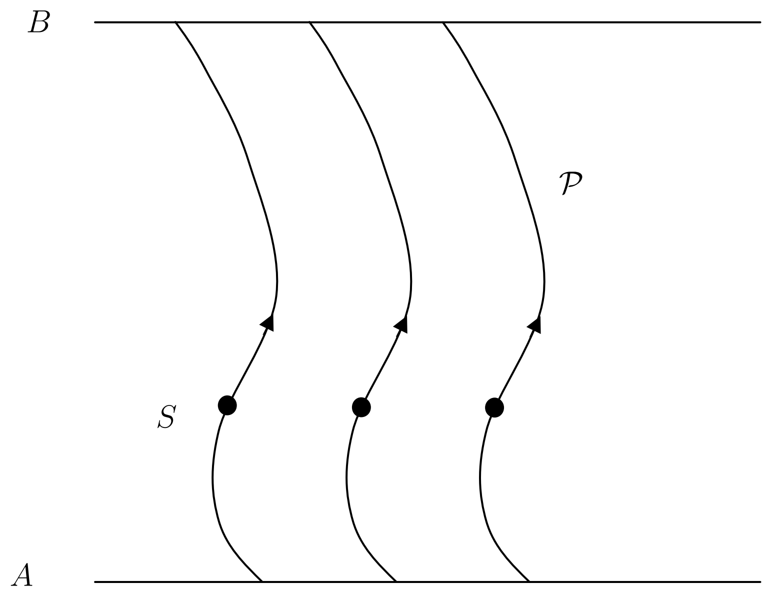

FIGURE 1. Illustration of Theorem 1.6

[13] proved their sufficiency for countable bipartite graphs. These criteria are more explicit, but in hindsight the concepts used in [27] are more fruitful.

Podewski and Steffens [28] made yet another important progress: they proved the conjecture for countable digraphs containing no infinite paths. Later, in [1], it was realized that this case can be easily reduced to the bipartite case, by the familiar device of doubling vertices in the digraph, thus transforming the digraph into a bipartite graph.

At that point in time there were two obstacles on the way to the proof of the conjecture - uncountability and the existence of infinite paths. The first of the two to be overcome was that of uncountability. In 1983 the marriage problem was solved for general cardinalities, in [11]. Soon thereafter, this was used to prove the infinite version of König's theorem [2]. Namely, the bipartite case of Theorem 1.6 was proved. Let us state it explicitly:

Theorem 1.7. *In any bipartite graph there exists a matching F and a cover C, such that C consists of the choice of precisely one vertex from each edge in F.*

As is well known, Hall's theorem fails in the infinite case. The standard example is that of the "playboy": take a graph with sides $M = \{m_0, m_1, m_2, \ldots\}$ and $W = \{w_1, w_2, \ldots\}$. For every $i > 0$ connect m_i to w_i, and connect m_0 (the playboy) to all w_i. Then every subset of M is connected to at least as many points in W as its size, and yet there is no marriage of M. This is just another indication that in the case of infinite matchings, cardinality is too crude a measure.

But Theorem 1.7 has an interesting corollary: that if "cardinality" is interpreted *in terms of the graph*, then Hall's theorem does apply also in the infinite case. Given two sets, I and J, of vertices in a graph G, we say that I is *matchable into* J if there exists an injection of I into J using edges of G. We write $I <_G J$ if I is matchable into J, but J is not matchable into I. (The ordinary notion of $|I| < |J|$ is obtained when G is the complete graph on a vertex set containing $I \cup J$.) A *marriage* of a side of a bipartite graph is a matching covering all its vertices. From Theorem 1.7 there follows:

Theorem 1.8. *Given a bipartite graph Γ with sides M and W, there does* not *exist a marriage of M if and only if there exists $A \subseteq M$, such that $N(A) <_\Gamma A$.*

To see how Theorem 1.8 follows from Theorem 1.7, assume that there is no marriage of M, and let F and C be as in Theorem 1.7. Let $I = M \setminus C$. Then the set of points connected to I is obviously $F[I]$ (the set of points connected by F to I), which is matchable by F into I. If there existed a matching K of I, then $K \cup (F \upharpoonright (M \cap C))$ would be a marriage of M, contrary to assumption. Thus I is unmatchable. The other implication in the theorem is obvious.

Proof-wise, the order is in fact reverse: Theorem 1.8 is proved first, and from it Theorem 1.7 follows, in a way that will be explained later, in Section 5.

By the result of [1], there follows from Theorem 1.7 also Theorem 1.6 for all graphs containing no infinite (unending or non-starting) paths. Thus there remained the problem of infinite paths. The difficulty they pose is that when one tries to "grow" the disjoint paths desired in the conjecture, they may end up being infinite, instead of being A–B-paths. In fact, in [1] it is proved that Theorem 1.6 is true, if one allows in $\mathcal{P}$ not only A–B-paths, but any paths that if they start at all, they do so at A, and if they end they do so at B.

The first breakthrough in the struggle against infinite paths was made in [3], where the countable case of the conjecture was proved. An equivalent, Hall-type, conjecture, was formulated, and the latter was proved for countable digraphs. The core of the proof was in a lemma, stating that if the Hall-like condition is satisfied, then any point in A can be linked to B by a path, whose removal leaves the Hall-like condition intact. The lemma is quite easy to prove in the bipartite case and also in graphs containing no unending paths, but in the general countable case it requires new tools and methods. Later, the sufficiency of the Hall-like condition for linkability (linking A into B by disjoint paths) was proved for graphs in which all but countably many points of A are linked to B [6], and Theorem 1.6 was proved for such graphs in [9].

In [8] a reduction was shown of the $\aleph_1$ case of the conjecture to the above mentioned lemma. Namely, a proof of the conjecture was given for digraphs of size $\aleph_1$, assuming that the lemma is true for such digraphs. Combined with a proof of the lemma for graphs with no unending paths, and for graphs with countable outdegrees, this settled the conjecture for digraphs of size at most $\aleph_1$, satisfying one of those properties. Optimistically, [8] declares that this reduction should probably work for general graphs.

The breakthrough leading to the solution of the general case was indeed the proof of this lemma for general graphs. As claimed in [8], the way from the lemma to the proof of the theorem indeed follows the same outline as in the $\aleph_1$ case. But the general case demands quite a bit more effort.

For the sake of relative self containment of the paper, most results from previous papers will be re-proved.

2. Notation

2.1. **Graph-theoretic notation.** One non-standard notation that we shall use is this: for a directed edge $e = (x, y)$ in a digraph we write $x = tail(e)$ and

$y = head(e)$. The rest of the notation is mostly standard, but here are a few reminders. Given a digraph D and a subset X of $V(D)$ we write $D[X]$ for the graph induced by D on X. Given a set U of vertices in an undirected graph, we denote by $N(U)$ the set of neighbors of vertices of U. In a digraph we write $N^+(U)$ (respectively $N^-(U)$) for the set of out-neighbors (respectively in-neighbors) of U. Adopting a common abuse of notation, when U consists of a single vertex u, we write $N(u), N^+(u), N^-(u)$ for $N(\{u\}), N^+(\{u\}), N^-(\{u\})$, respectively. Similar abuse of notation will apply also to other notions, without explicit mention.

2.2. **Webs.** A *web* Γ is a triple (D, A, B), where $D = D(\Gamma)$ is a digraph, and $A = A(\Gamma), B = B(\Gamma)$ are subsets of $V(D) = V(\Gamma)$. We usually write V for $V(D)$ and E for $E(D)$. If the identity of a web is not specified, we shall tacitly assume that the above notation - namely Γ, D, A and B - applies to it.

Assumption 2.1. Throughout the paper we shall assume that there are no edges going out of B, or into A.

Given a digraph D, we write $\overleftarrow{D}$ for the graph having the same vertex set as D, with all edges reversed. For a web $\Gamma = (D, A, B)$ we denote by $\overleftarrow{\Gamma}$ the web $(\overleftarrow{D}, B, A)$.

2.3. **Paths.** Following customary definitions, unless otherwise stated, a "path" in this paper is assumed to be simple, i.e. not self intersecting. *All paths P considered in the paper are assumed (unless empty) to have an initial vertex, denoted by $in(P)$.* If P is finite then its terminal vertex is denoted by $ter(P)$. The vertex set of a path P is denoted by $V(P)$, and its edge set by $E(P)$. The (possibly empty) path obtained by removing $in(P)$ and $ter(P)$ from P is denoted by P°.

Given a path P, we write $\overleftarrow{P}$ for the path in $\overleftarrow{D}$ obtained by traversing P in reverse order.

Given two vertices u, v on a path P, we write $u \le_P v$ (resp. $u <_P v$) if u precedes v on P (resp. u precedes v on P and $u \neq v$).

Given a set $\mathcal{P}$ of paths, we write $\mathcal{P}^f$ for the set of finite paths in $\mathcal{P}$, and $\mathcal{P}^\infty$ for the set of infinite paths in $\mathcal{P}$. We also write $V[\mathcal{P}] = \bigcup\{V(P) : P \in \mathcal{P}\}$, $E[\mathcal{P}] = \bigcup\{E(P) : P \in \mathcal{P}\}$, $in[\mathcal{P}] = \{in(P) : P \in \mathcal{P}\}$, and $ter[\mathcal{P}] = \{ter(P) : P \in \mathcal{P}^f\}$.

For a vertex x, we denote by (x) the path whose vertex set is $\{x\}$, having no edges.

For $X, Y \subseteq V$, a finite path P is said to be an *X–Y-path* if $in(P) \in X$ and $ter(P) \in Y$.

Given a path P and a vertex $v \in V(P)$, we write Pv for the part of P up to and including v, and vP for the part of P from v (including v) and on. If $Q = Pv$ for some $v \in V(P)$ we say that P is a *forward extension* of Q and write $P \vec{\succeq} Q$.

Given two paths, P and Q with $ter(P) = in(Q)$, we write $P * Q$, or sometimes just PQ, for the concatenation of P and Q, namely the (not necessarily simple) path, whose vertex set is $V(P) \cup V(Q)$ and whose edge set is $E(P) \cup E(Q)$. If $V(P) \cap V(Q) = \{ter(P)\} = \{in(Q)\}$ then $P * Q$ is a simple path. In this case clearly $P * Q \vec{\succeq} P$. Given paths P, Q sharing a common vertex v, we write PvQ for the (not necessarily simple) path $Pv * vQ$.

2.4. **Warps.** A set of vertex disjoint paths is called a *warp* (a term taken from weaving). If all paths in a warp are finite, then we say that the warp is of *finite* *character* (f.c.). A warp $\mathcal{W}$ is called X-*starting* if $in[\mathcal{W}] \subseteq X$. Given two sets of vertices, X and Y, a warp $\mathcal{W}$ is called an X–Y-warp if for every $P \in \mathcal{W}$ we have $in(P) \in X$, $ter(P) \in Y$ and $V(P) \cap (X \cup Y) = \{in(P), ter(P)\}$. We say that a warp $\mathcal{W}$ *links* X to Y if for every $x \in X$ there exists some $P \in \mathcal{W}$ such that $V(P) \cap X = \{x\}$ and $V(xP) \cap Y \neq \emptyset$. Note that a warp linking X to Y needs not be an X–Y warp, namely the initial points of its paths need not lie in X, and the terminal points do not necessarily lie in Y. An X–Y-warp linking X to Y is called an X–Y- *linkage*. An A–B-linkage in a web $\Gamma = (D, A, B)$ is called a *linkage of* Γ. A web having a linkage is called *linkable*. We write $\overleftarrow{\mathcal{W}}$ for the warp $\{\overleftarrow{P} \mid P \in \mathcal{W}\}$ in $\overleftarrow{D}$.

For a set $X \subseteq V$, we denote by $\langle X \rangle$ the warp consisting of all vertices of X as singleton paths. For every warp $\mathcal{W}$ we write $ISO(\mathcal{W})$ (standing for "isolated vertices of $\mathcal{W}$") for the set of vertices appearing in $\mathcal{W}$ as singleton paths.

Notation 2.2. Given a warp $\mathcal{W}$ and a set of vertices X, we write $\mathcal{W}[X]$ for the unique warp whose vertex set is $X \cap V[\mathcal{W}]$ and whose edge set is $\{(u, v) \in E[\mathcal{W}] \mid u, v \in X\}$. Paths in $\mathcal{W}[X]$ are sub-paths of paths in $\mathcal{W}$. Note that a path in $\mathcal{W}$ may break into more than one path in $\mathcal{W}[X]$. We also write $\mathcal{W} - X$ for $\mathcal{W}[V \setminus X]$.

Definition 2.3. A warp $\mathcal{U}$ is said to be an *extension* of a warp $\mathcal{W}$ if $V[\mathcal{W}] \subseteq V[\mathcal{U}]$ and $E[\mathcal{W}] \subseteq E[\mathcal{U}]$. We write then $\mathcal{W} \preccurlyeq \mathcal{U}$. Note that $\mathcal{U}$ may amalgamate paths in $\mathcal{W}$. If in addition $in[\mathcal{W}] = in[\mathcal{U}]$ then we say that $\mathcal{U}$ is a *forward extension* of $\mathcal{W}$ and write $\mathcal{U} \vec{\succcurlyeq} \mathcal{W}$. Note that in this case each path in $\mathcal{U}$ is a forward extension of some path in $\mathcal{W}$.

Notation 2.4. Given a warp $\mathcal{W}$ and a set $X \subseteq V$, we write $\mathcal{W}\langle X \rangle$ for the set of paths in $\mathcal{W}$ intersecting X, and $\mathcal{W}\langle \sim X \rangle$ for $\mathcal{W} \setminus \mathcal{W}\langle X \rangle$. Given two sets of vertices, X and Y, we write $\mathcal{W}\langle X, Y \rangle$ for $\mathcal{W}\langle X \rangle \cap \mathcal{W}\langle Y \rangle$ and $\mathcal{W}\langle X, \sim Y \rangle$ for $\mathcal{W}\langle X \rangle \cap \mathcal{W}\langle \sim Y \rangle$.

Given a vertex $x \in V[\mathcal{W}]$ we write $\mathcal{W}(x)$ for the path in $\mathcal{W}$ containing x (to be distinguished from $\mathcal{W}\langle x \rangle$, the set consisting of the single path $\mathcal{W}(x)$).

Given a warp $\mathcal{W}$ in a web (D, A, B), we write $\mathcal{W}_G$ for $\mathcal{W}\langle A \rangle$ and $\mathcal{W}_H$ for $\mathcal{W} \setminus \mathcal{W}_G$ (the subscript "G" stands for "ground" - these are the paths in $\mathcal{W}$ that start "from the ground", namely at A. The subscript "H" stands for "hanging in air". These terms originate in the way the authors are accustomed to draw webs - with the "A" side at the bottom, and the "B" side on top).

A set $\mathcal{F}$ of paths is called a *fractured warp* if its edge set is the edge set of a warp and every two paths $P, Q \in \mathcal{F}$ may intersect only if none of them is a trivial path and $in(P) = ter(Q)$ or $in(Q) = ter(P)$. If $\mathcal{W}$ is a warp and X is a set of vertices, we write $\mathcal{W} \downharpoonleft X$ for the fractured warp consisting of all paths of the form xPy where $P \in \mathcal{W}$, $x \in X \cup \{in(P)\}$, $y \in X \cup \{ter(P)\}$, $V(xPy) \not\subseteq X$ and $V(xPy) \cap X \subseteq \{x, y\}$. A somewhat more comprehensible definition is given by the following properties: $E[\mathcal{W} \downharpoonleft X] = E[\mathcal{W}] \setminus E[\mathcal{W}[X]]$, no path in $\mathcal{W} \downharpoonleft X$ amalgamates two paths in $\mathcal{W}$, and all singleton paths $(y) \in \mathcal{W}$, where $y \notin X$, belong to $\mathcal{W} \downharpoonleft X$.

A set of pairwise disjoint paths and directed cycles is called a *cyclowarp*. If $\mathcal{C}$ is a cyclowarp, we denote by $\mathcal{C}^{path}$ the warp obtained from $\mathcal{C}$ by removing all its cycles.

2.5. **Operations between warps.**

Notation 2.5. Let $\mathcal{U}$ and $\mathcal{W}$ be warps such that $V[\mathcal{U}] \cap V[\mathcal{W}] \subseteq ter[\mathcal{U}] \cap in[\mathcal{W}]$. Denote then by $\mathcal{U} * \mathcal{W}$ the warp $\{P * Q \mid\ P \in \mathcal{U}^f,\ Q \in \mathcal{W},\ in(Q) = ter(P)\} \cup \{P \in \mathcal{U} \mid ter(P) \notin in[\mathcal{W}]\}$. In particular, $\mathcal{U} * \mathcal{W}$ contains $\mathcal{U}^\infty$. Denote by $\mathcal{U} \diamond \mathcal{W}$ the warp whose vertex set is $V[\mathcal{U}] \cup V[\mathcal{W}]$ and whose edge set is $E[\mathcal{U}] \cup E[\mathcal{W}]$.

Thus $\mathcal{U} \diamond \mathcal{W} \supseteq \mathcal{U} * \mathcal{W}$. The difference is that $\mathcal{U} \diamond \mathcal{W}$ may contain also paths in $\mathcal{W}$ not meeting any path from $\mathcal{U}$.

There is also a binary operation defined on *all* pairs of warps. Given warps $\mathcal{U}$ and $\mathcal{W}$, their "arrow" $\mathcal{U}^\frown\mathcal{W}$ is obtained by taking each path in $\mathcal{U}$ and "carrying it along $\mathcal{W}$", if possible, and if not keeping it as it is. Formally, this is defined as follows:

Notation 2.6. Let $\mathcal{U}$ and $\mathcal{W}$ be two warps and let P be a path in $\mathcal{U}$. We define the *$\mathcal{U}$-$\mathcal{W}$-extension* $Ext_{\mathcal{U}-\mathcal{W}}(P)$ of P as follows. Consider first the case that P is finite. Let $u = ter(P)$. If there exists a path $Q \in \mathcal{W}$ satisfying $u \in V(Q)$ and $V(uQ) \cap V[\mathcal{U}] = \{u\}$ let $Ext_{\mathcal{U}-\mathcal{W}}(P) = PuQ$. In any other case (i.e. if either P is infinite or $u \notin V[\mathcal{W}]$ or $V(u\mathcal{W}(u))$ meets $\mathcal{U}$ at a vertex other than u) we take $Ext_{\mathcal{U}-\mathcal{W}}(P) = P$. Let

$$\mathcal{U}^\frown\mathcal{U} = \{Ext_{\mathcal{U}-\mathcal{W}}(P) :\ P \in \mathcal{U}\}.$$

(See Figure 2.)

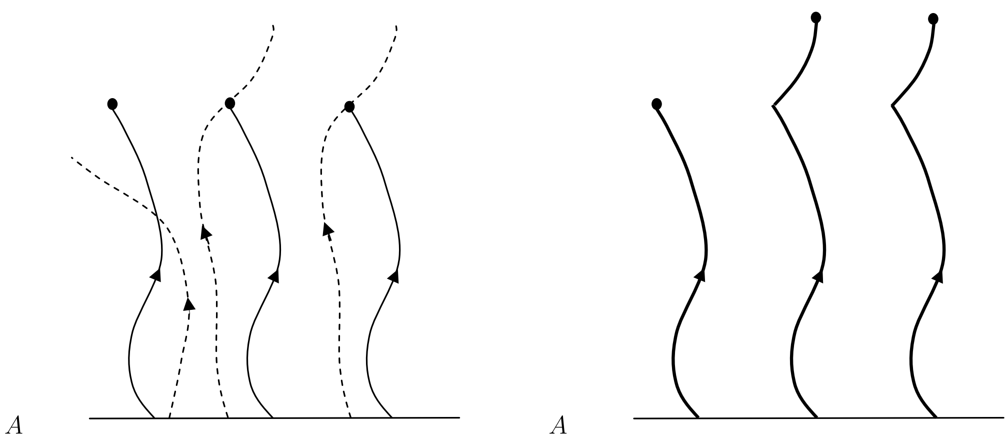

FIGURE 2. On the left there are drawn a warp $\mathcal{U}$ (solid line) and a warp $\mathcal{W}$ (dashed line). On the right is drawn their "arrow", $\mathcal{U}^\frown\mathcal{W}$.

Note that $\mathcal{U}^\frown\mathcal{W}$ is a warp and $\mathcal{U}^\frown\mathcal{W} \overrightarrow{\geq} \mathcal{U}$.

Observation 2.7. *$\mathcal{W} \overrightarrow{\geq} \mathcal{U}$ if and only if $\mathcal{U}^\frown\mathcal{W} = \mathcal{W}$.*

Next we wish to define the "arrow" of a sequence of warps. As a first step, we define the limit of an ordinal-indexed sequence of warps.

Definition 2.8. Let $(S_\alpha :\ \alpha < \theta)$ be a sequence of sets. The *limit* (actually, $\liminf$) of the sequence, denoted by $\lim_{\alpha<\theta} S_\alpha$, is defined as $\bigcup_{\beta<\theta} \bigcap_{\beta \leq \alpha < \theta} S_\alpha$. Let $(\mathcal{W}_\alpha :\ \alpha < \theta)$ be a sequence of warps. The *limit* $\lim_{\alpha<\theta} \mathcal{W}_\alpha$ of the sequence is the warp whose edge set is $\lim_{\alpha<\theta} E[\mathcal{W}_\alpha]$ and whose vertex set is $\lim_{\alpha<\theta} V[\mathcal{W}_\alpha]$.

As noted, $\lim_{\alpha<\theta}\mathcal{W}_\alpha$ is in fact the "lim inf" of the warps. The fact that it is indeed a warp is straightforward. Note that by this definition if θ is not a limit ordinal, namely $\theta = \psi + 1$, then $\lim_{\alpha<\theta}\mathcal{W}_\alpha$ is just $\mathcal{W}_\psi$.

Observation 2.9. *Let $(\mathcal{W}_\alpha : \alpha < \theta)$ be a sequence of warps. Then $ter[\lim_{\alpha<\theta}\mathcal{W}_\alpha] \supseteq \lim_{\alpha<\theta} ter[\mathcal{W}_\alpha]$.*

Definition 2.10. Let $(\mathcal{W}_\alpha : \alpha < \theta)$ be an ordinal-indexed sequence of warps. Define a sequence $\mathcal{W}'_\alpha,\ \alpha < \theta$, by: $\mathcal{W}'_0 = \mathcal{W}_0$, $\mathcal{W}'_{\psi+1} = \mathcal{W}'_\psi \frown \mathcal{W}_{\psi+1}$ (where $\psi + 1 < \theta$), and for limit ordinals $\alpha \leq \theta$ define $\mathcal{W}'_\alpha = \lim_{\psi<\alpha}\mathcal{W}'_\psi$. Let $\uparrow_{\alpha<\theta}\mathcal{W}_\alpha$ be defined as $\mathcal{W}'_\theta$ if θ is a limit ordinal, and as $\mathcal{W}'_\beta$ if $\theta = \beta + 1$.

Note that if $(\mathcal{W}_\alpha : \alpha < \theta)$ is $\vec{\lesssim}$-ascending, then this definition coincides with the "limit" definition. If $\{\mathcal{W}_i,\ i \in I\}$ is an unordered set of warps, then $\uparrow_{i\in I}\mathcal{W}_i$ can be defined by first imposing an arbitrary well-order on I. Of course, the resulting warp depends on the order chosen, but when applied we shall use a fixed well order.

2.6. **Almost disjoint families of paths.** Given a set X of vertices, a set $\mathcal{P}$ of paths is called X-*joined* if the intersection of the vertex sets of any two paths from $\mathcal{P}$ is contained in X (so, a warp is just a $\emptyset$-joined family of paths). For a single vertex x, we write simply "x-joined" instead of "$\{x\}$-joined". A family of x-joined paths starting at x is called a *fan*. A family of x-joined paths terminating at x is called an *in-fan*.

Given a set $X \subseteq V$ and a vertex u, a u-fan $\mathcal{F}$ is said to be a u–X-fan if $ter[\mathcal{F}] \subseteq X$. An X–u-in-fan is defined similarly. A u-fan consisting of infinite paths is called a (u,∞)-fan.

2.7. **Separation.**

Definition 2.11. An A–B-*separating* set of vertices in a web $\Gamma = (D, A, B)$ is plainly said to be *separating*.

Definition 2.12. Given a (not necessarily separating) subset S of $V(D)$, a vertex $s \in S$ is said to be *essential* (for separation) in S if it is not separated from B by $S \setminus \{s\}$. The set of essential elements of S is denoted by $\mathcal{E}(S)$, and the set $S \setminus \mathcal{E}(S)$ of inessential vertices by $\mathcal{IE}(S)$. If $S = \mathcal{E}(S)$ then we say that S is *trimmed*.

Convention 2.13. By removing those vertices of A from which B is unreachable, we may assume that A is trimmed. We shall tacitly make this assumption.

Lemma 2.14. *If S is an A–B separating set of vertices, then so is $\mathcal{E}(S)$.*

Proof. Let Q be an A–B-path. Since by assumption S is A–B separating, $V(Q) \cap S \neq \emptyset$. The last vertex s on Q belonging to S is essential in S, since the path sQ shows that s is not separated from B by $S \setminus \{s\}$. □

A path P in a warp $\mathcal{W}$ is said to be *essential* (in $\mathcal{W}$) if P is finite and $ter(P) \in \mathcal{E}(ter[\mathcal{W}])$. The set of essential paths in $\mathcal{W}$ is denoted by $\mathcal{E}(\mathcal{W})$, and the set of inessential paths by $\mathcal{IE}(\mathcal{W})$. If $\mathcal{W} = \mathcal{E}(\mathcal{W})$ we say that $\mathcal{W}$ is *trimmed*.

To Definition 1.3 we add the following. Given a set X of vertices, a vertex set S is called X-∞-*separating* if it contains a vertex on every infinite path starting in X. The minimal size of an X-∞-separating set is denoted by $\sigma(X,\infty)$.

Notation 2.15. For a set S of vertices in a web $\Gamma = (D, A, B)$ we denote by $RF(S) = RF_\Gamma(S)$ the set of all vertices separated by S from B. We also write $RF^\circ(S) = RF(S) \setminus \mathcal{E}(S)$.

The letters "RF" stand for "roofed", a term originating again in the way the authors draw their webs, with the "A" side at the bottom, and the "B" above. Note that in particular, $S \subseteq RF(S)$ and $\mathcal{IE}(S) \subseteq RF^\circ(S)$. Given a warp $\mathcal{W}$, we write $RF(\mathcal{W}) = RF(ter[\mathcal{W}])$, $RF^\circ(\mathcal{W}) = RF^\circ(ter[\mathcal{W}])$. A warp $\mathcal{W}$ is said to be *roofed* by a set of vertices S if $V[\mathcal{W}] \subseteq RF(S)$.

Lemma 2.16. *Let S be a set of vertices and P any path. If $V(P) \cap RF(S) \neq \emptyset$ then the last vertex on P belonging to $RF(S)$ belongs to $\mathcal{E}(S) \cup \{ter(P)\}$.*

Proof. Let v be the last vertex on P belonging to $RF(S)$. Suppose that $v \neq ter(P)$. We have to show that $v \in \mathcal{E}(S)$. Let u be the vertex following v on P. Then $u \notin RF(S)$, meaning that there exists an S-avoiding path Q from u to B. Since $v \in RF(S)$ the path vuQ meets $\mathcal{E}(S)$. Since this meeting can occur only at v, it follows that $v \in \mathcal{E}(S)$. □

Lemma 2.17. *If C, D are sets of vertices such that $\mathcal{E}(D) \subseteq C \subseteq D$ then $\mathcal{E}(C) = \mathcal{E}(D)$.*

Proof. Let $x \in \mathcal{E}(D)$. Then there exists an x–B path avoiding $D \setminus \{x\}$, and thus avoiding $C \setminus \{x\}$, showing that $x \in \mathcal{E}(C)$. On the other hand, if $x \in \mathcal{E}(C)$ then there exists an x–B path P avoiding $C \setminus \{x\}$. If P does not avoid $D \setminus \{x\}$ then its last vertex belongs to $\mathcal{E}(D)$, and thus to C, a contradiction. Thus $x \in \mathcal{E}(D)$. □

Observation 2.18. *Let S, T, X, Y be four sets of vertices, with $X \cap Y = \emptyset$. If $X \subseteq RF(T \cup Y)$ and $Y \subseteq RF(S \cup X)$ then $X \cup Y \subseteq RF(S \cup T)$ (otherwise stated as: $\mathcal{E}(S \cup T \cup X \cup Y) = \mathcal{E}(S \cup T)$).*

Proof. For an $(X \cup Y)$–B path P consider the last vertex z on P belonging to $X \cup Y$. By the conditions of the observation, zP must meet $S \cup T$. □

Lemma 2.19. *If R, S, T are three sets of vertices satisfying $T = \mathcal{E}(T)$ and $RF(R) \subseteq RF(S) \subseteq RF(T)$ then S is R–T-separating.*

Proof. Consider an R–T path P and let $x = ter(P)$. Since $T = \mathcal{E}(T)$ there exists an x-B path Q satisfying $in(Q) = x$ and $V(Q) \cap T = \{x\}$. Then PxQ is an R–B path. (A-priory, PxQ might not be a simple path. However, it obviously *contains* a simple R–B path.) Since S is R-B separating, we have $V(PxQ) \cap S \neq \emptyset$. But since $S \subseteq RF(T)$ and $V(Q) \cap T = \{x\}$, we have $V(Q) \cap RF(S) \subseteq \{x\}$, and hence $V(PxQ) \cap S = V(P) \cap S \neq \emptyset$, proving the lemma. □

Notation 2.20. Let S be a set of vertices in a web $\Gamma = (D, A, B)$, such that $RF(S) = S$ (which is equivalent to S being equal to $RF(T)$ for some set T). We denote then by $\Gamma[S]$ the web $(D[S], S \cap A, \mathcal{E}(S))$. Given a warp $\mathcal{W}$ we write $\Gamma[\mathcal{W}]$ for $\Gamma[RF(\mathcal{W})]$.

Lemma 2.21. *Let $(S_\alpha : \alpha < \theta)$ be a sequence of sets, satisfying $S_\alpha \subseteq RF(S_\beta)$ for $\alpha < \beta < \theta$. Then $RF(\lim_{\alpha<\theta} S_\alpha) \supseteq \bigcup_{\alpha<\theta} RF(S_\alpha)$.*

Proof. Let $x \in \bigcup_{\alpha<\theta} RF(S_\alpha)$. We may assume that $x \in RF(S_0)$ and thus $x \in \bigcap_{\alpha<\theta} RF(S_\alpha)$. Let P be an x–B path and let t be the last vertex on P belonging to $\bigcup_{\alpha<\theta} S_\alpha$. Say, $t \in S_\beta$. Since $t \in RF(S_\alpha)$ for all $\beta < \alpha < \theta$ the vertex t must be in S_α for all such α, and hence $t \in \lim_{\alpha<\theta} S_\alpha$. □

2.8. **Deletion and quotient.** A basic operation on webs is that of removing vertices. In fact, there are two ways of doing this. One is plain deletion: for a subset X of V we denote by $\Gamma - X$ the web $(D - X, A \setminus X,\ B \setminus X)$. For a path P we abbreviate and write $\Gamma - P$ instead of $\Gamma - V(P)$.

Lemma 2.22. *$RF(X \cup Y) = X \cup RF_{\Gamma - X}(Y)$.*

Proof. Note that X is contained in the sets appearing on both sides of the equality. Hence it suffices to show that for a vertex v not belonging to X, namely a vertex of $\Gamma - X$, a v–B path avoids $X \cup Y$ in Γ if and only if it avoids Y in $\Gamma - X$. But this is almost a tautology. □

The other type of removal is taking a quotient. The difference from deletion is that taking a quotient with respect to a set X of vertices means deleting the vertices of X as vertices through which paths can go from A to B, but also adding X to A, indicating a commitment to link X to B. If indeed such linking is possible, then the possibility arises of linking vertices of A to B by first linking them to X, and then linking vertices of X to B.

Definition 2.23. Given a subset X of $V \setminus A$, write D/X for the digraph obtained from D by deleting all edges going into vertices of X, and all vertices in $RF^\circ(X)$, including those of $\mathcal{IE}(X)$. Define Γ/X as the web $(D/X, \mathcal{E}(A \cup X), B)$.

Observation 2.24. *Since we are assuming that A is trimmed, $A(\Gamma/X) = (A \cup X) \setminus RF^\circ(X)$.*

Remark 2.25. In bipartite webs deleting a vertex $b \in B$ and taking a quotient with respect to it are the same, as far as linkability is concerned, since taking a quotient with respect to b means that b is added to A, and is linked automatically to itself. This is the reason why the quotient operation is not needed in the proof of the bipartite case of the theorem.

Lemma 2.26. *For any two sets X and Y of vertices, $RF^\circ_\Gamma(X \cup Y) = RF^\circ(X) \cup RF_{\Gamma/X}(Y \setminus RF^\circ(X))$.*

Proof. Let v be a vertex in $RF^\circ_\Gamma(X \cup Y)$. Suppose that $v \notin RF^\circ(X) \cup RF_{\Gamma/X}(Y \setminus RF^\circ(X))$. Then there exists a path P from v to B in Γ/X, avoiding $(Y \setminus RF^\circ(X)) \setminus \{v\}$. Since P is contained in $V(\Gamma/X)$, it is disjoint from $RF^\circ(X)$. Hence the fact that it avoids $(Y \setminus RF^\circ(X)) \setminus \{v\}$ means that in fact it avoids $Y \setminus \{v\}$, and since there are no edges in Γ/X going into X, it also avoids $(X \cup Y) \setminus \{v\}$, contradicting the assumption on v. This proves that the set on the left hand side is contained in the set on the right hand side. The other containment relation is proved similarly. □

Lemma 2.27. *For any two sets X and Y of vertices, if $Y \cap RF^\circ(X) = \emptyset$ then $\Gamma/(X \cup Y) = (\Gamma/X)/Y$.*

Proof. Let us first show that the two webs share the same vertex set. By the definition of the quotient, we need to show:

$$V \setminus RF^\circ_\Gamma(X \cup Y) = ((V \setminus RF^\circ_\Gamma(X)) \setminus RF^\circ_{\Gamma/X}(Y).$$

This follows from Lemma 2.26.

We also need to show that the two webs share the same source set, namely that $\mathcal{E}_\Gamma(A \cup X \cup Y) = \mathcal{E}_{\Gamma/X}(\mathcal{E}_\Gamma(A \cup X) \cup Y)$. But this also follows from Lemma 2.26.

It remains to show equality of the edge sets of the two webs. Write $V' = V(\Gamma/(X \cup Y)) = V((\Gamma/X)/Y)$ and $A' = A(\Gamma/(X \cup Y)) = A((\Gamma/X)/Y)$. Then $E(\Gamma/(X \cup Y)) = E((\Gamma/X)/Y) = \{(u,v) \in E(\Gamma) \mid\ u \in V',\ v \in V' \setminus A'\}$. □

Corollary 2.28. *For any two sets X_1 and X_2 of vertices, if $Y = \mathcal{E}(X_1 \cup X_2)$ then $(\Gamma/X_1)/Y = (\Gamma/X_2)/Y = \Gamma/Y$.*

Given a warp $\mathcal{W}$, we write $\Gamma/\mathcal{W}$ for $\Gamma/ter[\mathcal{W}]$. If $\mathcal{U}$ and $\mathcal{W}$ are two warps, we write $\mathcal{U}/\mathcal{W}$ for $\mathcal{U}/ter[\mathcal{W}]$.

Definition 2.29. Given a warp $\mathcal{W}$ and a set X of vertices, we define the quotient $\mathcal{W}/X$ by $V[\mathcal{W}/X] = (V[\mathcal{W}] \cup X) \setminus RF^\circ(X)$ and $E[\mathcal{W}/X] = \{(u,v) \in E[\mathcal{W}] \mid\ u \notin RF^\circ(X), v \notin RF(X)\}$.

We end this section with a few lemmas, some of which are obvious and some have similar proofs to those above, and hence we list them without proofs:

Lemma 2.30. *$\mathcal{W}/X$ is a warp in Γ/X.*

Lemma 2.31. *$\langle \mathcal{E}(X) \setminus V[\mathcal{W}] \rangle \subseteq \mathcal{W}/X$.*

Lemma 2.32. *If $in[\mathcal{W}] \subseteq A(\Gamma)$ then $in[\mathcal{W}/X] \subseteq A(\Gamma/X)$*

Lemma 2.33. *If $\mathcal{W} \preccurlyeq \mathcal{W}'$ then $\mathcal{W}/X \preccurlyeq \mathcal{W}'/X$. If $\mathcal{W} \vec{\preccurlyeq} \mathcal{W}'$ then $\mathcal{W}/X \vec{\preccurlyeq} \mathcal{W}'/X$.*

Lemma 2.34. *$in[\mathcal{W}/X] = (in[\mathcal{W}] \cup X) \setminus RF^\circ(X)$ and $ter[\mathcal{W}/X] \supseteq (ter[\mathcal{W}] \setminus RF^\circ(X)) \cup (\mathcal{E}(X) \setminus V[\mathcal{W}])$.*

Lemma 2.35. *For a subset Z of $V(\Gamma)$ and a warp $\mathcal{V}$ in Γ we have $RF^\circ_\Gamma(\mathcal{V}) \cap V(\Gamma/Z) \subseteq RF^\circ_{\Gamma/Z}(\mathcal{V}/Z)$.*

Lemma 2.36. *If S, T are disjoint sets of vertices, then $RF_{\Gamma - T}(S) \setminus RF^\circ(T) \subseteq RF_{\Gamma/T}(S \setminus RF^\circ(T))$.*

3. Waves and hindrances

Definition 3.1. An A-starting warp $\mathcal{W}$ is called a *wave* if $ter[\mathcal{W}]$ is A–B-separating.

Clearly, $\langle A \rangle$ (namely, the set of singleton paths, $\{(a) \mid\ a \in A\}$), is a wave. It is called *the trivial wave*.

Observation 3.2. *If $S = RF(S) \supseteq A$ and $\mathcal{W}$ is a wave in $\Gamma[S]$ then $\mathcal{W}$ is also a wave in Γ.*

Lemma 2.14 implies:

Lemma 3.3. *If $\mathcal{W}$ is a wave then so is $\mathcal{E}(\mathcal{W})$.*

This gives

Lemma 3.4. *A path W belonging to a wave $\mathcal{W}$ is essential in $\mathcal{W}$ if and only if $\mathcal{W} \setminus \{W\}$ is not a wave.*

Proof. If W is inessential, then by Lemma 3.3, $A \subseteq RF(\mathcal{E}(\mathcal{W})) \subseteq RF(\mathcal{W} \setminus \{W\})$.

If, on the other hand, W is essential, then W is finite. Let $t = ter(W)$. Since $t \in \mathcal{E}(ter[\mathcal{W}])$, there exists a path P from t to B avoiding $ter[\mathcal{W}] \setminus \{t\}$, and then WtP is an A–B path avoiding $ter[\mathcal{W} \setminus \{W\}]$, showing that $\mathcal{W} \setminus \{W\}$ is not a wave. □

One nice property of waves is that they stay waves upon taking quotients.

Lemma 3.5. *If $\mathcal{U}$ is a wave and $X \subseteq V$ then $\mathcal{U}/X$ is a wave in Γ/X.*

Proof. By Lemmas 2.30 and 2.32, The warp $\mathcal{U}/X$ is indeed an $A(\Gamma/X)$-starting warp in Γ/X.

Let Q be a path in Γ/X from $A(\Gamma/X)$, namely $(A \cup X) \setminus RF^{\circ}(X)$, to B. We have to show that Q meets $ter[\mathcal{U}/X]$.

If $in(Q) \in A$ then, since $\mathcal{U}$ is a wave, $in(Q) \in RF_{\Gamma}[\mathcal{U}]$. Otherwise $in(Q) \in \mathcal{E}(X)$. Thus in both cases $in(Q) \in RF_{\Gamma}[\mathcal{U}] \cup \mathcal{E}(X)$. Let t be the last vertex on Q belonging to $RF_{\Gamma}[\mathcal{U}] \cup \mathcal{E}(X)$. From the choice of t it follows that $t \notin RF^{\circ}_{\Gamma}(X) \cup RF^{\circ}_{\Gamma}(\mathcal{U})$, and hence $t \in (ter[\mathcal{U}] \setminus RF^{\circ}_{\Gamma}(X)) \cup (\mathcal{E}(X) \setminus RF^{\circ}_{\Gamma}(\mathcal{U}))$. By Lemma 2.34 $t \in ter[\mathcal{U}/X]$. ☐

A wave $\mathcal{W}$ is called a *hindrance* if $in[\mathcal{W}] \neq A$. The origin of the name is that in finite webs a hindrance is an obstruction for linkability. In the infinite case this is not necessarily so. A web containing a hindrance is said to be *hindered*.

As a corollary of Lemma 3.5 we have

Corollary 3.6. *If $A \cap RF(S) = \emptyset$ and $\mathcal{H}$ is a hindrance in Γ then $\mathcal{H}/S$ is a hindrance in Γ/S.*

For, if $a \in A \setminus in[\mathcal{H}]$ then $a \in A(\Gamma/S) \setminus in[\mathcal{H}/S]$.

Clearly, a hindrance is a non-trivial wave. A web not containing any non-trivial wave is called *loose*.

Lemma 3.7 (the self roofing lemma)**.** *If $\mathcal{W}$ is a wave then $V[\mathcal{W}] \subseteq RF(\mathcal{W})$.*

Proof. Suppose, for contradiction, that there exists a path Q avoiding $ter[\mathcal{W}]$, from some vertex x on a path $P \in \mathcal{W}$ to B. Taking a sub-path of Q, if necessary, we can assume that PxQ is a path. Then PxQ avoids $ter[\mathcal{W}]$, contradicting the fact that $\mathcal{W}$ is a wave. ☐

Corollary 3.8. *Let $X \subseteq V$ and let $\mathcal{W}$ be a wave in $\Gamma - X$. Then $V[\mathcal{W}] \setminus ter[\mathcal{W}] \subseteq RF^{\circ}(ter[\mathcal{W}] \cup X)$*

Proof. Let $u \in V[\mathcal{W}] \setminus ter[\mathcal{W}]$. By Lemma 3.7 we have $V[\mathcal{W}] \subseteq RF_{\Gamma - X}(\mathcal{W}) \subseteq RF_{\Gamma}(ter[\mathcal{W}] \cup X)$. Since $u \notin ter[\mathcal{W}] \cup X$, we get $u \in RF^{\circ}(ter[\mathcal{W}] \cup X)$. ☐

Definition 3.9. A warp $\mathcal{W}$ is called *self roofing* if $V[\mathcal{W}] \subseteq RF(\mathcal{W})$.

Lemma 3.7 implies that every wave is self roofing. In fact, an easy corollary of this lemma together with Lemma 3.5 extends it to waves in quotient webs.

Corollary 3.10. *If $\mathcal{W}$ is a wave in Γ/X for some set X then $\mathcal{W}$ is a self roofing warp in Γ.*

For two waves $\mathcal{W}$ and $\mathcal{W}'$ we write $\mathcal{W} \equiv \mathcal{W}'$ if $ter[\mathcal{E}(\mathcal{W})] = ter[\mathcal{E}(\mathcal{W}')]$. Also write $\mathcal{W} \leq \mathcal{U}$ if $RF(\mathcal{W}) \subseteq RF(\mathcal{U})$. Clearly, this is equivalent to the statement that $ter[\mathcal{W}] \subseteq RF(\mathcal{U})$. The relation $\leq$ is a partial order on the equivalence classes of the $\equiv$ relation. Namely, if $\mathcal{W} \leq \mathcal{U}$ and $\mathcal{W} \equiv \mathcal{W}'$, $\mathcal{U} \equiv \mathcal{U}'$ then $\mathcal{W}' \leq \mathcal{U}'$, while if $\mathcal{W} \leq \mathcal{U}$ and $\mathcal{U} \leq \mathcal{W}$ then $\mathcal{U} \equiv \mathcal{W}$. We write $\mathcal{U} > \mathcal{W}$ if $\mathcal{W} \leq \mathcal{U}$ and $\mathcal{W} \not\equiv \mathcal{U}$, i.e., $RF(\mathcal{W}) \subsetneq RF(\mathcal{U})$. We say that a wave $\mathcal{W}$ is $\leq$-*maximal* if there is no wave $\mathcal{U}$ satisfying $\mathcal{U} > \mathcal{W}$.

By the self roofing lemma (Lemma 3.7) we have:

Corollary 3.11. *For two waves $\mathcal{U}$ and $\mathcal{W}$, if $\mathcal{W} \preccurlyeq \mathcal{U}$ then $\mathcal{W} \leq \mathcal{U}$.*

The next lemma is formulated in great generality (hence its complicated statement), so as to avoid repeating the same type of arguments again and again:

Lemma 3.12. *Let X and Y be two sets of vertices in Γ, and let $\mathcal{U}, \mathcal{W}$ be warps, satisfying the following conditions:*
(1) $\mathcal{U}$ is a wave in $\Gamma - X$.
(2) $Y \subseteq RF_{\Gamma - X}(\mathcal{U})$.
(3) $\mathcal{W}$ is a self roofing warp in $\Gamma - Y$.
(4) $X \subseteq RF_{\Gamma - Y}(\mathcal{W})$ and $X \cap V[\mathcal{W}] \subseteq in[\mathcal{W}]$.
(5) Every path in $\mathcal{W}$ meets $RF_{\Gamma - X}(\mathcal{U})$.
Then $\mathcal{E}_\Gamma(ter[\mathcal{U}^\frown\mathcal{W}]) = \mathcal{E}_\Gamma(ter[\mathcal{U}] \cup ter[\mathcal{W}]) = \mathcal{E}_\Gamma(ter[\mathcal{U}] \cup ter[\mathcal{W}] \cup X \cup Y)$.

(The last equality means of course that $X \cup Y \subseteq RF(ter[\mathcal{U}] \cup ter[\mathcal{W}])$.)

Proof. By (1) and (2) we have $Y \subseteq RF(X \cup ter[\mathcal{U}])$ and by (3) and (4) we have $X \subseteq RF(Y \cup ter[\mathcal{W}])$. This together with Observation 2.18 yields $\mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}]) = \mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}] \cup X \cup Y)$, so we only need to show the first equality. Since $ter[\mathcal{U}^\frown\mathcal{W}] \subseteq ter[\mathcal{U}] \cup ter[\mathcal{W}]$, by Lemma 2.17 it suffices to show that $ter[\mathcal{U}^\frown\mathcal{W}] \supseteq \mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}])$.

Let $z \in \mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}])$. We need to show that $z \in ter[\mathcal{U}^\frown\mathcal{W}]$.

Consider first the case that $z \in ter[\mathcal{U}]$. If $z \notin V[\mathcal{W}]$ then $\mathcal{U}(z) \in \mathcal{U}^\frown\mathcal{W}$ and we are done. Thus we may assume that $z \in V[\mathcal{W}]$, which by (3) entails that $z \in RF(ter[\mathcal{W}] \cup Y)$ and $z \notin Y$. The fact that $z \in \mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}] \cup X \cup Y)$ implies therefore that $z \in ter[\mathcal{W}]$, again implying $\mathcal{U}(z) \in \mathcal{U}^\frown\mathcal{W}$.

We are left with the case that $z \in ter[\mathcal{W}] \setminus ter[\mathcal{U}]$. Let $W = \mathcal{W}(z)$ and let u be the last vertex in W which is in $RF_{\Gamma - X}(\mathcal{U})$. The existence of such u is certified by (5). Note that by (4) we know that the path uW does not meet X. Therefore we may apply Lemma 2.16 in the web $\Gamma - X$ and get $u \in ter[\mathcal{U}] \cup \{z\}$. Suppose that $u = z$. Then by by the choice of u and by (2), we have $z \notin X \cup Y$. Since $z \in \mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}] \cup X \cup Y)$, there exists a z–B path avoiding $ter[\mathcal{U}] \cup ter[\mathcal{W}]$. By the choice of u and by (1), this path must meet $ter[\mathcal{U}]$, and the only vertex at which this can happen is u itself, contradicting the assumption of the present case. We have thus proved that $u \neq z$, and thus $u \in ter[\mathcal{U}]$. This implies that $\mathcal{U}(u)uW \in \mathcal{U}^\frown\mathcal{W}$, proving $z \in ter[\mathcal{U}^\frown\mathcal{W}]$.

□

The most frequently used case of this lemma will be that of $Y = X = \emptyset$:

Lemma 3.13. *If $\mathcal{U}$ and $\mathcal{W}$ are waves then so is $\mathcal{U}^\frown\mathcal{W}$.*

Proof. Combine the lemma with the fact that $ter[\mathcal{U}]$, and hence *a fortiori* $ter[\mathcal{U}] \cup ter[\mathcal{W}]$, is A–B-separating. □

Another case we shall use is in which $X = \emptyset$ but Y is not necessarily empty.

Lemma 3.14. *If $\mathcal{U}$ is a wave in Γ, $Y \subseteq RF(\mathcal{U})$ and $\mathcal{W}$ is a wave in $\Gamma - Y$, then $\mathcal{U}^\frown\mathcal{W}$ is a wave in Γ.*

Taking $Y = \emptyset$ but X not necessarily empty, and using Corollary 3.10 we get:

Lemma 3.15. *Let X, Z be subsets of $V(\Gamma)$ such that $X \subseteq Z$. Let $\mathcal{U}$ be a wave in $\Gamma - X$ and let $\mathcal{W}$ be a wave in Γ / Z. If every path in $\mathcal{W}$ meets $RF_{\Gamma - X}(\mathcal{U})$ then $\mathcal{U}^\frown\mathcal{W}$ is a wave in Γ.*

Remark 3.16. Since $\mathcal{U}$ is a wave in $\Gamma - X$, every path in $\mathcal{W}$ starting at A meets $RF_{\Gamma-X}(\mathcal{U})$. Therefore the only paths in $\mathcal{W}$ for which the assertion of meeting $RF_{\Gamma-X}(\mathcal{U})$ really needs to be checked are those starting at Z.

By Corollary 3.11 if $\mathcal{U}$ and $\mathcal{W}$ are waves, then $\mathcal{U} \leq \mathcal{U}\frown\mathcal{W}$. Lemma 3.12 implies more:

Lemma 3.17. *For any two waves $\mathcal{U}$ and $\mathcal{W}$ we have: $\mathcal{U}, \mathcal{W} \leq \mathcal{U}\frown\mathcal{W}$.*

Lemma 3.18. *$\mathcal{E}(ter[\mathcal{U}\frown\mathcal{W}]) \cap RF^\circ(\mathcal{U}) = \emptyset$.*

Proof. $\mathcal{E}(ter[\mathcal{U}\frown\mathcal{W}]) \cap RF^\circ(\mathcal{U}) \subseteq \mathcal{E}(ter[\mathcal{U}] \cup ter[\mathcal{W}]) \cap RF^\circ(ter[\mathcal{U}] \cup ter[\mathcal{W}]) = \emptyset$ □

Lemma 3.19. *If $(\mathcal{W}_\alpha : \alpha < \theta)$ is a $\vec{\preccurlyeq}$-ascending sequence of waves, then $\uparrow_{\alpha<\theta} \mathcal{W}_\alpha$ is a wave and $\uparrow_{\alpha<\theta} \mathcal{W}_\alpha \geq \mathcal{W}_\alpha$ for every $\alpha < \theta$.*

Proof. This is a direct corollary of Observation 2.9 and Lemma 2.21. □

Since clearly $\uparrow_{\alpha<\theta} \mathcal{W}_\alpha \vec{\succcurlyeq} \mathcal{W}_\alpha$ for all $\alpha < \theta$, by Zorn's lemma this implies:

Lemma 3.20. *In every web there exists a $\vec{\preccurlyeq}$-maximal wave. Furthermore, every wave can be forward extended to a $\vec{\preccurlyeq}$-maximal wave.*

One corollary of this lemma is that a hindered web contains a maximal hindrance.

Corollary 3.21. *If there exists in Γ a hindrance then there exists in Γ a $\vec{\preccurlyeq}$-maximal wave that is a hindrance.*

Next we show that there is no real distinction between $\vec{\preccurlyeq}$-maximality and $\leq$-maximality.

Lemma 3.22. *Any $\vec{\preccurlyeq}$-maximal wave (and hence also any $\preccurlyeq$-maximal wave) is $\leq$-maximal. If $\mathcal{V}$ is a $\leq$-maximal wave then there does not exist a trimmed wave $\mathcal{W}$ such that $\mathcal{E}(\mathcal{V}) \lneqq \mathcal{W}$.*

Proof. Assume first that $\mathcal{V}$ is a $\leq$-non-maximal wave, i.e., there exists a wave $\mathcal{W} > \mathcal{V}$, meaning that $RF(\mathcal{W}) \supsetneqq RF(\mathcal{V})$. By Lemma 3.17 it follows that $\mathcal{V}\frown\mathcal{W} \neq \mathcal{V}$, and since $\mathcal{V}\frown\mathcal{W} \vec{\succcurlyeq} \mathcal{V}$ it follows that $\mathcal{V}$ is not $\vec{\preccurlyeq}$-maximal and hence also not $\preccurlyeq$-maximal. This proves the first part of the lemma.

Assume next that $\mathcal{V}$ is a $\leq$-maximal wave. Let $\mathcal{U} = \mathcal{E}(\mathcal{V})$. Suppose, for contradiction, that $\mathcal{U} \lneqq \mathcal{W}$ for some trimmed wave $\mathcal{W}$. This means that there exists some path $W \in \mathcal{W} \setminus \mathcal{U}$. Since $\mathcal{W}$ is trimmed, W is finite. Write $t = ter(W)$. Since $in[W] \subseteq A$ and we assume no edges enter A, the only two possibilities are that either W is a proper forward extension of some path in $\mathcal{U}$ or W does not meet $V[\mathcal{U}]$ at all. In both case we have $t \notin ter[\mathcal{U}]$. Since $\mathcal{W}$ is trimmed we have $t \notin RF^\circ(\mathcal{W})$ and hence $t \notin RF^\circ(\mathcal{U})$. Thus $t \notin RF(\mathcal{U})$, which implies that $\mathcal{W} > \mathcal{V}$, a contradiction. □

Corollary 3.23. *If $\mathcal{U}, \mathcal{V}$ are each either $\preccurlyeq$-maximal, or $\vec{\preccurlyeq}$-maximal, or $\leq$-maximal waves, then $\mathcal{U} \equiv \mathcal{V}$.*

Proof. By the lemma, in all cases $\mathcal{U}$ and $\mathcal{V}$ are $\leq$-maximal. By Lemma 3.17 $\mathcal{U}\frown\mathcal{V} \geq \mathcal{U}, \mathcal{V}$, which, by the $\leq$-maximality of $\mathcal{U}$ and $\mathcal{V}$, implies that $RF(\mathcal{U}\frown\mathcal{V}) = RF(\mathcal{U}) = RF(\mathcal{V})$. The last equality means that $\mathcal{U} \equiv \mathcal{V}$. □

Thanks to Corollary 3.23, we may speak about "maximal waves", without specifying whether we mean $\leq$ or $\preccurlyeq$ or $\vec{\preccurlyeq}$-maximality, as long as do this only in contexts involving vertices roofed by the waves, or quotient over the waves, or other properties that do not distinguish between equivalent waves.

Observation 3.24. *If $\mathcal{W}$ is a wave, then $A(\Gamma/\mathcal{W}) = \mathcal{E}(ter[\mathcal{W}])$.*

Proof. Recall that $\Gamma/\mathcal{W}$ is defined as $\Gamma/ter[\mathcal{W}]$, which in turn means that $A(\Gamma/\mathcal{W}) = (A \cup ter[\mathcal{W}]) \setminus RF^\circ(ter[\mathcal{W}])$. Since $\mathcal{E}(ter[\mathcal{W}]) = ter[\mathcal{W}] \setminus RF^\circ(ter[\mathcal{W}])$ we have $\mathcal{E}(ter[\mathcal{W}]) \subseteq (A \cup ter[\mathcal{W}]) \setminus RF^\circ(ter[\mathcal{W}])$. Since $\mathcal{W}$ is a wave, $A \subseteq RF(\mathcal{W})$, implying that $A \setminus RF^\circ(\mathcal{W}) \subseteq ter[\mathcal{W}]$, and hence $(A \cup ter[\mathcal{W}]) \setminus RF^\circ(ter[\mathcal{W}]) \subseteq ter[\mathcal{W}] \setminus RF^\circ(ter[\mathcal{W}]) = \mathcal{E}(ter[\mathcal{W}])$. □

Lemma 3.25. *If $\mathcal{W}$ is a wave in Γ and $\mathcal{V}$ is a wave in $\Gamma/\mathcal{W}$ then $\mathcal{W} * \mathcal{V}$ is a wave in Γ.*

Proof. Let P be a path from A to B. We have to show that P meets $ter[\mathcal{W} * \mathcal{V}]$. Since $\mathcal{W}$ is a wave, P meets $ter[\mathcal{W}]$. Let t be the last vertex on P belonging to $ter[\mathcal{W}]$. Then clearly $t \in \mathcal{E}(ter[\mathcal{W}])$ and $V(tP) \cap RF^\circ(\mathcal{W}) = \emptyset$, and hence by Observation 3.24 tP is an $A(\Gamma/\mathcal{W})$–$B(\Gamma/\mathcal{W})$ path in $\Gamma/\mathcal{W}$. Thus tP meets $ter[\mathcal{V}]$, and since clearly $ter[\mathcal{V}] \subseteq ter[\mathcal{W} * \mathcal{V}]$ it follows that tP meets $ter[\mathcal{W} * \mathcal{V}]$ and so does P, as required. □

Lemma 3.26. *If $\mathcal{W}$ is a $\preccurlyeq$-maximal wave then $\Gamma/\mathcal{W}$ is loose.*

Proof. Assume, for contradiction, that there exists a non-trivial wave $\mathcal{V}$ in $\Gamma/\mathcal{W} = \Gamma/\mathcal{E}(\mathcal{W})$. If all paths in $\mathcal{V}$ are singletons then, since $\mathcal{V}$ is non-trivial, $\mathcal{V} \subsetneqq \langle ter[\mathcal{E}(\mathcal{W})]\rangle$, contradicting the definition of $\mathcal{E}(\mathcal{W})$. Thus not all paths in $\mathcal{V}$ are singletons, and hence $\mathcal{W} * \mathcal{V} \supsetneqq \mathcal{W}$, and since by Lemma 3.25 $\mathcal{W} * \mathcal{V}$ is a wave this contradicts the maximality of $\mathcal{W}$. □

By Lemma 3.23, the $\preccurlyeq$-maximality in the above lemma can be replaced by $\vec{\preccurlyeq}$- or $\leq$-maximality.

Lemma 3.27. *Let X be a subset of $V \setminus A$, and let $\mathcal{U}$ be a warp in Γ avoiding X, such that $\mathcal{U}$ is a wave in $\Gamma - X$. Then $\mathcal{U}/X$ is a wave in Γ/X. Furthermore,*

$$RF_{\Gamma - X}(\mathcal{U}) \setminus RF^\circ(X) \subseteq RF_{\Gamma/X}(\mathcal{U}/X). \tag{1}$$

Proof. Note that $\langle \mathcal{E}(X)\rangle \subseteq \mathcal{U}/X$. Since $A(\Gamma/X) \subseteq (RF_{\Gamma-X}(\mathcal{U}) \setminus RF^\circ(X)) \cup \mathcal{E}(X)$, in order to prove that $\mathcal{U}/X$ is a wave in Γ/X it suffices to prove (1). Let Q be a path in Γ/X starting at a vertex $z \in RF_{\Gamma-X}(\mathcal{U}) \setminus RF^\circ(X)$ and ending in B. We have to show that Q meets $ter[\mathcal{U}/X]$. If Q meets X then it meets $\mathcal{E}(X)$ and we are done. If not, then the desired conclusion follows from the fact that $z \in RF_{\Gamma-X}(\mathcal{U})$. □

A corollary of this lemma is that Γ/X contains more "advanced" waves than $\Gamma - X$:

Corollary 3.28. *If X and $\mathcal{U}$ are as above, and if $\mathcal{V}$ is a maximal wave in Γ/X, then $RF_\Gamma(\mathcal{V}) \supseteq RF_{\Gamma-X}(\mathcal{U})$ and $RF^\circ_\Gamma(\mathcal{V}) \supseteq RF^\circ_{\Gamma-X}(\mathcal{U})$.*

One advantage that the quotient operation has over deletion is the following. Given two sets of vertices, X_1 and X_2, there is no natural way of combining a wave in $\Gamma - X_1$ with a wave in $\Gamma - X_2$, so as to yield a third wave in some web. By contrast, there does exist a natural definition of a combination of a wave $\mathcal{W}_1$ in Γ/X_1 with a wave $\mathcal{W}_2$ in Γ/X_2. Writing $X = \mathcal{E}(X_1 \cup X_2)$, we can combine $\mathcal{W}_1$ and $\mathcal{W}_2$ by taking the warp $(\mathcal{W}_1/X)^\frown(\mathcal{W}_2/X)$.

Lemma 3.29. *Let $X_1, X_2 \subseteq V$, and write $X = \mathcal{E}(X_1 \cup X_2)$. If $\mathcal{W}_1$ is a wave in Γ/X_1 and $\mathcal{W}_2$ is a wave in Γ/X_2, then $(\mathcal{W}_1/X)^\frown(\mathcal{W}_2/X)$ is a wave in Γ/X. Moreover,*

$$RF_{\Gamma/X}((\mathcal{W}_1/X)^\frown(\mathcal{W}_2/X)) \supseteq RF_{\Gamma/X}(\mathcal{W}_1/X) \cup RF_{\Gamma/X}(\mathcal{W}_2/X).$$

Proof. Corollary 2.28 and Lemma 3.5 imply that $\mathcal{W}_1/X$ and $\mathcal{W}_2/X$ are both waves in Γ/X, and hence by Lemma 3.13 so is $(\mathcal{W}_1/X)^\frown(\mathcal{W}_2/X)$. The second part of the lemma follows from Lemma 3.17. □

The next lemma is a special case of Lemma 3.19 that we will need.

Lemma 3.30. *Let $(X_i \ : \ 0 \leq i < \omega)$ be a $\subseteq$-ascending sequence of subsets of $V \setminus A$. For each $i < \omega$, let $\mathcal{W}_i$ be a wave in Γ/X_i. Write $X = \mathcal{E}(\bigcup_{i<\omega} X_i)$. Then $\uparrow_{i<\omega} (\mathcal{W}_i/X)$ (taken as an up-arrow of waves in Γ/X) is a wave in Γ/X.*

We conclude this section with two lemmas taken from [3], whose proofs are rather technical and hence will not be presented here:

Lemma 3.31. *If Γ is hindered and X is a finite subset of $V \setminus A$ then $\Gamma - X$ is hindered.*

This is not necessarily true if X is infinite.

Lemma 3.32. *If Γ is unhindered, and $\Gamma - v$ is hindered for a vertex $v \in V \setminus A$, then there exists a wave $\mathcal{W}$ in Γ such that $v \in ter[\mathcal{W}]$.*

4. Bipartite conversion of webs and warp-alternating paths

4.1. **Aims of this section.** As already mentioned, Menger's theorem is better understood, in both its finite and infinite cases, if its relationship to König's theorem is apparent. As mentioned in the introductin, a simple transformation, observed in [1] (but probably known earlier), reduces the finite case of Menger's theorem to König's theorem. This "bipartite conversion" is effective also for webs containing no infinite paths, but not for general webs. We chose to describe it here since it inspired many of the ideas of the present proof, and some points in the proof are illuminated by it. The bipartite conversion is also the most natural source for definitions involving alternating paths. As is common in matching theory, the latter will constitute one of our main tools.

4.2. **The bipartite conversion of a web.** The "bipartite conversion" turns a digraph into a bipartite graph. Every vertex of the digraph is replaced by two copies, one sending arrows and the other receiving them. The graph becomes then bipartite, with one side consisting of the "sending" copies, and the other consisting of the "receiving" copies.

For webs the construction is a little different: A-vertices are given only "sending" copies, and B-vertices are given only "receiving" copies. Thus the web $\Gamma = (G, A, B)$ turns into a bipartite web $\Delta = \Delta(\Gamma) = (G_\Delta, A_\Delta, B_\Delta)$, in the following way. Every vertex $v \in V \setminus A$ is assigned a vertex $w(v) \in B_\Delta$, and every vertex $v \in V \setminus B$ is assigned a vertex $m(v) \in A_\Delta$. Thus, vertices in $V \setminus (A \cup B)$ are assigned two copies each. The edge set $E_\Delta = E(G_\Delta)$ is defined as $\{(m(x), w(y)) \mid (x, y) \in E(G)\} \cup \{(m(x), w(x)) \mid x \in V \setminus (A \cup B)\}$.

The above transformation converts a web into a bipartite web, together with a matching, namely the set of edges $\{(m(x), w(x)) \mid x \in V \setminus (A \cup B)\}$. This transformation can be reversed: given a bipartite graph Δ whose two sides are A and B, together with a matching J in it, one can construct from it a web $\Lambda = \Lambda(J)$ (the reference to Δ is suppressed), as follows. To every edge $(x, y) \in J$ we assign a vertex $v(x, y)$. The vertex set $V(\Lambda)$ is $\{v(x, y) \mid (x, y) \in J\} \cup V(\Delta) \setminus \bigcup J$. (Here $\bigcup J$ is the set of vertices participating in edges from J.) The "source" side A_Λ of Λ is defined as $A_\Delta \setminus \bigcup J$, and the "destination" set B_Λ is $B_\Delta \setminus \bigcup J$.

For $u \in V(\Lambda)$ define $m(u) = u$ if $u \in A_\Lambda \setminus J$, and $m(v(x, y)) = x$ (namely, the A-vertex of (x, y)) for every edge $(x, y) \in J$. Let $w(u) = u$ if $u \in B_\Lambda \setminus J$, and $w(v(x, y)) = y$ (namely, the B-vertex of (x, y)) for every edge $(x, y) \in J$. The edge set of Λ is defined as $\{(u, v) \mid (m(u), w(v)) \in E[\Delta]\}$.

Let us now return to our web Γ, and consider a warp $\mathcal{W}$ in it. Let $J = J(\mathcal{W})$ be the matching in $\Delta(\Gamma)$, defined by $J = \{(m(u), w(v)) \mid (u, v) \in E[\mathcal{W}]\} \cup \{(m(u), w(u)) \mid u \notin \bigcup E[\mathcal{W}]\}$. We abbreviate and write $\Lambda(\mathcal{W})$ for $\Lambda(J(\mathcal{W}))$. From the definitions there easily follows:

Lemma 4.1. *If $\mathcal{W}$ is a linkage in Γ, then $J(\mathcal{W})$ is a marriage of A_Δ in $\Delta = \Delta(\Gamma)$. If Γ does not contain unending paths, then the converse is also true.*

4.3. **Alternating paths.** The definition of one of our main tools, that of $\mathcal{Y}$-alternating paths, where $\mathcal{Y}$ is a warp, is quite involved. To be able to follow its fine points, it may be helpful to keep in mind the main property required of a $\mathcal{Y}$-alternating path: that the symmetric difference of its edge set and the edge set of $\mathcal{Y}$ is the edge set of a warp. For a precise definition, see Definition 4.3 below.

Definition 4.2. Let $\mathcal{Y}$ be a warp in Γ. A *$\mathcal{Y}$-alternating path* is a sequence Q having one of the following forms:

(i) an infinite sequence $(u_0, F_0, w_1, R_1, u_1, F_1, w_2, R_2, u_2, \ldots)$,
(ii) an infinite sequence $(w_1, R_1, u_1, F_1, w_2, R_2, u_2, \ldots)$,
(iii) $(u_0, F_0, w_1, R_1, u_1, F_1, w_2, R_2, u_2, \ldots, R_k, u_k)$,
(iv) $(u_0, F_0, w_1, R_1, u_1, F_1, w_2, R_2, u_2, \ldots, R_k, w_k, F_{k+1}, w_{k+1})$,
(v) $(w_1, R_1, u_1, F_1, w_2, R_2, u_2, \ldots, R_k, u_k)$,
(vi) $(w_1, R_1, u_1, F_1, w_2, R_2, u_2, \ldots, R_k, w_k, F_{k+1}, w_{k+1})$,

and satisfying the following conditions:

(1) u_i, w_i are vertices, and F_i, R_i are paths having at least one edge each. Furthermore, $in(F_i) = u_i$, $ter(F_i) = w_{i+1}$ for all relevant values of i. The paths R_i are subpaths of paths from $\mathcal{Y}$, and $in(R_i) = u_i$, $ter(R_i) = w_i$ for all relevant values of i.
(2) For paths of types (i),(iii) or (iv) $u_0 \notin V[\mathcal{Y}]$.
(3) For paths of types (iv) or (vi) $w_{k+1} \notin V[\mathcal{Y}]$.
(4) If $v \in V(R_i) \cap V(R_j)$ for $i \neq j$, then either $v = u_i = w_j$ or $v = w_i = u_j$.
(5) If $v \in V(F_i) \cap V(F_j)$ for $i \neq j$, then either $v = u_i = w_{j+1}$ or $v = w_{i+1} = u_j$.
(6) If $V(F_i) \cap V(R_j) \neq \emptyset$ then either:
  (i) $j = i + 1$, and $V(F_i) \cap V(R_j) = \{w_j\}$, or:
  (ii) $i > j$ and $V(F_i) \cap V(R_j) \cap \{u_i, w_{i+1}, u_j, w_j\} = \emptyset$, namely the paths F_i and R_j meet only at their interiors, or:

(iii) $j = i$, and $w_i, w_{i+1} \notin V(F_i) \cap V(R_j)$, namely the paths F_i and R_j meet only at u_i and possibly also at their interiors.

The notation "F_i" and "R_i" stands for "forward" and "reverse", respectively - we think of Q as going forward on F_i, and reversely on R_i. The links F_i and R_i are called "forward links" and "backward links" of Q, respectively. The last three requirements in the definition mean that links can only meet at their endpoints, with one exception: a forward link can go through an internal vertex of a backward link, if the latter precedes it in the path. Allowing this may seem redundant, since if this happens then the alternating path can be replaced by a shorter one having the same initial and terminal vertices. But there is one place, namely Lemma 4.13 below, in which this type of paths must be permitted.

The first vertex (u_0 or w_1) on Q is denoted by $in(Q)$. If $in(Q) = u_0 \in A$ then Q is said to be A-*starting*. Note that by condition (2) this implies that $u_0 \notin V[\mathcal{Y}]$. If Q is infinite, then Q is said to be an $(in(Q), \infty)$-$\mathcal{Y}$- alternating path. If Q is finite, we write $ter(Q)$ for its last vertex, and say that Q is an $(in(Q), ter(Q))$-$\mathcal{Y}$-alternating path. If $in(Q) = u_0 \in A \setminus V[\mathcal{Y}]$ and $ter(Q) \in B \setminus V[\mathcal{Y}]$, we say that Q is *augmenting*. The source of this name is that in this case the application of Q to $\mathcal{Y}$ adds one more path to $\mathcal{Y}$ than it removes from it (see Definition 4.3 below for the meaning of "application" of an alternating path to a warp). This meaning of "augmentation" does not depend on the condition $in(Q) \in A,\ ter(Q) \in B$, but since this is the only case we shall use the notion of augmentation, we add this condition.

If $in(Q) = w_1 \in ter[\mathcal{Y}]$ and $ter(Q) = u_k \in in[\mathcal{Y}]$ then Q is said to be *reducing*. In this case the application of Q to $\mathcal{Y}$ removes one more path from $\mathcal{Y}$ than it adds to it.

If Q is infinite, or it is finite and $ter(Q) \notin V[\mathcal{Y}]$ then Q is said to be $\mathcal{Y}$-*leaving*.

Definition 4.3. For a $\mathcal{Y}$-alternating path Q as above, $\mathcal{Y} \triangle Q$ is the cyclowarp whose edge set is $E[\mathcal{Y}] \triangle E(Q)$, namely $E[\mathcal{Y}] \setminus \bigcup E(R_i) \cup \bigcup E(F_i)$, with $ISO(\mathcal{Y} \triangle Q) = ISO(\mathcal{Y})$.

(Recall that $ISO(\mathcal{Y})$ denotes the set of singleton paths in $\mathcal{Y}$.) The cyclowarp $\mathcal{Y} \triangle Q$ is also said to be the result of *applying* Q *to* $\mathcal{Y}$.

Definition 4.4. Let $\mathcal{U}, \mathcal{Y}$ be warps. A $\mathcal{Y}$-alternating path is said to be $[\mathcal{U}, \mathcal{Y}]$-alternating if all paths F_i in Definition 4.2 are subpaths of paths in $\mathcal{U}$. A $[\mathcal{U}, \mathcal{Y}]$-alternating path is said to be $\mathcal{U}$-*comitted* if no R_i contains a point from $V[\mathcal{U}] \setminus ter[\mathcal{U}]$ as an internal point. Namely, if the alternating path switches to $\mathcal{U}$ whenever possible.

Every $\mathcal{Y}$-alternating path in Γ corresponds in a natural way to a $J(\mathcal{Y})$-alternating path in $\Delta(\Gamma)$, which, in turn, corresponds to a path in $\Lambda(\mathcal{Y})$. Moreover, an augmenting $\mathcal{Y}$-alternating path corresponds to an A_Λ–B_Λ path in Λ. We summarize this in:

Lemma 4.5. *Let $\mathcal{Y}$ be a warp in Γ, and let $\Lambda = \Lambda(\mathcal{Y})$. Then there exists an augmenting $\mathcal{Y}$-alternating path if and only if there exists an A_Λ–B_Λ path in Λ.*

An A–B-warp $\mathcal{Y}$ is called *strongly maximal* if $|\mathcal{Y} \setminus \mathcal{U}| \geq |\mathcal{U} \setminus \mathcal{Y}|$ for every A–B-warp $\mathcal{U}$. The following is well known (see, e.g., [20]):

Lemma 4.6. *An A–B-warp $\mathcal{Y}$ is strongly maximal if and only if there does not exist an augmenting $\mathcal{Y}$-alternating path.*

Note that in the finite case "strong maximality" means just "having maximal size", and hence obviously there exist strongly maximal warps. Hence the following result implies Menger's theorem:

Theorem 4.7. *Let $\mathcal{Y}$ be a strongly maximal A–B-warp. For every $P \in \mathcal{Y}$ let $bl(P)$ be the last vertex on P participating in an A-starting $\mathcal{Y}$-alternating path if such a vertex exists, and $bl(P) = in(P)$ if there is no A-starting $\mathcal{Y}$-alternating path meeting P. Then the set $BL = \{bl(P) : P \in \mathcal{Y}\}$ is A–B-separating.*

(The letters "bl" stand for "blocking".) This result also yields an equivalent formulation of Theorem 1.6, noted in [20]: in every web there exists a strongly maximal A–B-warp.

Theorem 4.7 was proved by Gallai [15]. A detailed proof is given in Chapter 3 of [14]. We give here an outline of the proof, since it yields one of the simplest proofs of the finite case of Menger's theorem, and since the idea will recur in Section 8.

Proof of Theorem 4.7. Let T be an A–B-path. Let P be the first path from $\mathcal{Y}$ it meets, say at a vertex z. Assuming that $z \neq bl(P)$, it must precede $bl(P)$ on P, since it lies on the alternating path Tz. Assuming that T avoids BL, it follows that either:

(i) T meets a path $R \in \mathcal{Y}$ at a vertex $u \in V(R)$ preceding $bl(R)$ on R, and $uT - u$ is disjoint from $V[\mathcal{Y}]$, or:

(ii) T meets a path $R \in \mathcal{Y}$ at a vertex $u \in V(R)$ preceding $bl(R)$ on R, and the next vertex w on T belonging to $V(W)$ for some $W \in \mathcal{Y}$ comes after $bl(W)$ on$W$.

Assume that (i) is true. Let Z be a $\mathcal{G}$-alternating path from $bl(R)$ to $Y \setminus S$. If Z does not meet T, then $Tu\overleftarrow{R}bl(R)Z$ is an augmenting $\mathcal{G}$-alternating path, contradicting Lemma 4.6. If Z meets T, let z be the last vertex on Z belonging to $V(T)$. Then the path TzZ is again an augmenting $\mathcal{G}$-alternating path, again yielding a contradiction.

On the other hand, (ii) is impossible since the alternating path reaching $bl(R)$ can be extended by adding to it $\overleftarrow{R}uTw$, so as to form an alternating path meeting W beyond $bl(W)$. □

4.4. **Safe alternating paths.**

Definition 4.8. A $\mathcal{Y}$-alternating path Q is called *safe* if:

1. For every $P \in \mathcal{Y}$ the intersection $E[Q] \cap E(P)$ (which, in the notation of Definition 4.2, is $\bigcup E(R_i) \cap E(P)$) is the edge set of a subpath (that is, a single interval) of P, and:
2. $E(Q) \setminus E[\mathcal{Y}]$ does not contain an infinite path or a cycle.

We use the abbreviation "$\mathcal{Y}$-s.a.p" for "safe $\mathcal{Y}$-alternating path". A $\mathcal{Y}$-s.a.p whose forward links F_i are fragments of a warp $\mathcal{W}$ is called a $[\mathcal{W}, \mathcal{Y}]$-s.a.p.

If Q is an infinite $\mathcal{Y}$-alternating path then $\mathcal{Y} \triangle Q$ may contain infinite paths, even if $\mathcal{Y}$ itself is f.c (reminder - "f.c." means "of finite character", namely having no infinite paths). See Figure 3.

The name "safe" originates in the fact that this cannot occur if Q is safe. For, each path in $\mathcal{Y} \triangle Q$ consists then of only three parts (one or two of which may be empty) - a subpath of a path of $\mathcal{Y}$, followed by a path lying outside $\mathcal{Y}$, followed then by another subpath of a path of $\mathcal{Y}$. For the same reason, $\mathcal{Y} \triangle Q$ does not contain cycles. We summarize this in:

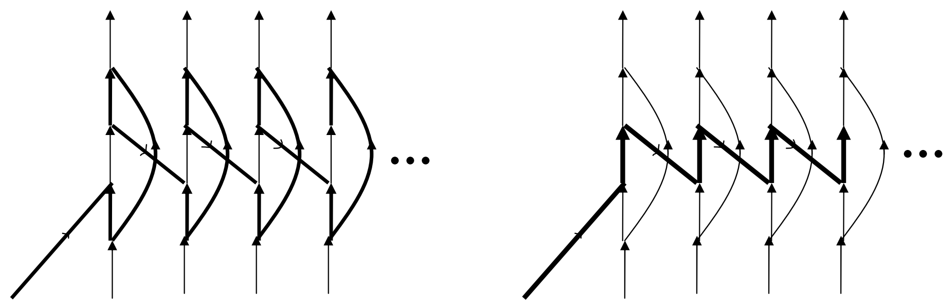

FIGURE 3. An example of an alternating path (bolded on the left) whose application results in a warp including an infinite path (bolded on the right).

Lemma 4.9. *If $\mathcal{Y}$ is warp of f.c. and Q is a $\mathcal{Y}$-s.a.p, then also $\mathcal{Y}\triangle Q$ is a warp of f.c.*

Definition 4.10. A (u,v)-$\mathcal{Y}$-alternating path Q (where possibly $v=\infty$) is called *degenerate* if $\mathcal{Y}\triangle Q$ contains a path from u to v.

The definition of "safeness" implies:

Lemma 4.11. *If a (u,v)-$[\mathcal{W},\mathcal{Y}]$-s.a.p Q is degenerate, then the path connecting u to v in $\mathcal{Y}\triangle Q$ is contained in a path from $\mathcal{W}$.*

A fact that we shall use about s.a.p's is:

Theorem 4.12. *Let $\mathcal{Z}$ and $\mathcal{Y}$ be f.c. warps, such that $in[\mathcal{Z}] \supseteq in[\mathcal{Y}]$. Then there exists a choice of a z-starting $\mathcal{Y}$-leaving maximal s.a.p $Q(z)$ for each $z \in in[\mathcal{Z}] \setminus in[\mathcal{Y}]$, such that those s.a.p's $Q(z)$ that are finite end at distinct vertices of $ter[\mathcal{Z}]$ (namely, $ter(Q(z)) \neq ter(Q(z'))$ whenever $z \neq z'$ and $Q(z), Q(z')$ are finite. Note that the paths $Q(z)$ themselves are not required to be disjoint).*

The maximality of the paths $Q(z)$ means that each $Q(z)$ is continued either indefinitely or until a vertex of $ter[\mathcal{Z}] \setminus V[\mathcal{Y}]$ is reached. For the proof of the theorem we shall need the following lemma:

Lemma 4.13. *Let $\mathcal{Z}$ and $\mathcal{Y}$ be f.c warps such that $in[\mathcal{Z}] \supseteq in[\mathcal{Y}]$, and let $u \in in[\mathcal{Z}] \setminus V[\mathcal{Y}]$. Then at least one of the following possibilities occurs:*

1. *There exists a (u,∞)-$[\mathcal{Z},\mathcal{Y}]$-s.a.p, or:*
2. *There exists a vertex $v \in ter[\mathcal{Z}] \setminus V[\mathcal{Y}]$ for which there exist both a (u,v)-$[\mathcal{Z},\mathcal{Y}]$-s.a.p and a (v,u)-$[\mathcal{Y},\mathcal{Z}]$ alternating path.*

Note that in case (2), the (u,v)-$[\mathcal{Z},\mathcal{Y}]$-s.a.p must be of type (iv) and the (v,u)-$[\mathcal{Y},\mathcal{Z}]$ alternating path must be reducing of type (v). To follow the logic of the proof, keep in mind that $[\mathcal{Y},\mathcal{Z}]$-alternating paths are $\mathcal{Z}$-alternating, but not necessarily $\mathcal{Y}$-alternating. Namely, they are "$\mathcal{Z}$-committed", meaning that whenever they meet a $\mathcal{Z}$-path they must switch to it, but they are not "$\mathcal{Y}$-committed". In contrast, $[\mathcal{Z},\mathcal{Y}]$-alternating paths are $\mathcal{Y}$-committed, while not necessarily $\mathcal{Z}$-committed. The following two examples illustrate this point:

Example 4.14. Suppose that $\mathcal{Y}$ consists of one path, $Y=(a,b,c,d)$, while $\mathcal{Z}$ consists of the paths $(a,d),(s,b,t)$ and (x,c,y). Consider first a case in which $u=x$. Since the graph is finite, (2) is impossible, and hence (1) should hold. Indeed, the easiest way to show that this is true is to take $v = t$. The safe $(u, v)-[\mathcal{Z}, \mathcal{Y}]$-alternating path (written by the order of its vertices) whould then be (x, c, b, t). Since this alternating path is $\mathcal{Z}$-committed, we can choose its reverse (t, b, c, x) to be the (v, u)-$[\mathcal{Y}, \mathcal{Z}]$-alternating path required in (1). Note, however, that the choice of $v = t$ is not unique. We could also take $v = y$, with the safe $[\mathcal{Z}, \mathcal{Y}]$-$(u, v)$-alternating path (x, c, b, a, d, c, y) and the (v, u)-$[\mathcal{Y}, \mathcal{Z}]$-alternating path (y, c, x).

Consider next another case, in which $u = s$. In this case, if we try to construct a $\mathcal{Z}$-committed alternating path, we end up with the alternating path (s, b, a, d, c, y), which is not safe. The only way to obtain (1) is then taking $v = t$. The safe (u, v)-$[\mathcal{Z}, \mathcal{Y}]$-alternating path is (s, b, a, d, c, b, t), and the $[\mathcal{Y}, \mathcal{Z}]$-alternating path is (t, b, s).

Proof of the lemma: By duplicating edges when necessary we may assume that $E[\mathcal{Z}] \cap E[\mathcal{Y}] = \emptyset$. It is clear that vertices on paths from $\mathcal{Y}$ not belonging to $V[\mathcal{Z}]$ do not play any role in the proof. They can be ignored, meaning that subpaths having them as internal points can be made to be single edges, and then terminal points of the resulting warp not belonging to $V[\mathcal{Z}]$ can be removed. Hence we shall assume that $V[\mathcal{Y}] \subseteq V[\mathcal{Z}]$.

Let SR (standing for "safely reachable") be the set of vertices $v \in ter[\mathcal{Z}] \setminus V[\mathcal{Y}]$ for which there exists a (u, v)-$[\mathcal{Z}, \mathcal{Y}]$-s.a.p, and let C be the set of vertices x for which there exists a (v, x)-$[\mathcal{Y}, \mathcal{Z}]$- alternating path $T(x)$ for some $v \in SR$. Note that $T(x)$ is not necessarily unique, but to avoid cumbersome phrasing we shall sometimes pretend that it is. Thus we shall refer by $T(x)$ to *some* alternating path satisfying the above conditions.

Assuming negation of possibility (2) of the lemma, we have $u \notin C$. Our aim is to show that this implies possibility (1) of the lemma. To that end, we construct a u-starting $[\mathcal{Z}, \mathcal{Y}]$-s.a.p S. This is done in stages, where at the i-th stage we shall have at hand a u-starting $[\mathcal{Z}, \mathcal{Y}]$-s.a.p S_i extending S_{i-1}, whose last link is a backward link on some path $Y_i \in \mathcal{Y}$, ending at a vertex $u_i \notin C$ (the paths Y_i need not be distinct). The construction will be shown to go on indefinitely, meaning that possibility (1) of the lemma is true.

Let $u_0 = u$. The path $Z_1 = \mathcal{Z}(u)$ must meet some path in $\mathcal{Y}$, or else $ter(Z_1) \in SR$, meaning that Z_1 can serve as $T(u)$ to show that $u \in C$. Let w_1 be the first vertex on Z_1 lying on some path $Y_1 \in \mathcal{Y}$. Since $u \notin V[\mathcal{Y}]$, we know that $w_1 \neq u$. Since $w_1 \notin in[\mathcal{Z}]$ and $in[\mathcal{Y}] \subseteq in[\mathcal{Z}]$ we also know that $w_1 \neq in(Y_1)$. Let y be the vertex preceding w_1 on Y_1.

Assertion 4.15. *Neither w_1 nor y are in C.*

If $w_1 \in C$ then extending $T(w_1)$ by the path $u\mathcal{Z}(u)$ (either as an extension of a link, in case the last link on $T(w_1)$ is a $\mathcal{Z}$-link, or as a separate link, if the last link on $T(w_1)$ is a $\mathcal{Y}$-link,) would show that $u \in C$. Similarly, if $y \in C$ then $T(y)$ can be extended by the forward link yY_1w_1 followed by the backward link $u\mathcal{Z}(u)$ to show that $u \in C$.

Returning to the construction of S, we go back on Y_1 to the first vertex u_1 on Y_1 not belonging to C (possibly $u_1 = in(Y_1)$). By the assertion, this means going at least one edge back on Y_1, meaning that the path S_1 obtained is $\mathcal{Y}$-alternating.

Let $B_1 = V((u_1Y_1w_1)^\circ)$ - this is the set of vertices on which S_1 goes backwards. (Recall that P° is obtained from a path P by removing from it $in(P)$ and $ter(P)$.)

Assume now that $i \geq 1$, and that $S_i = (u_0, F_1, w_1, R_1, \ldots R_i, u_i)$ has been already defined, where each forward link F_i is part of a path $Z_i = \mathcal{Z}(u_{i-1})$, the backward link R_i is part of a path $Y_i = \mathcal{Y}(w_i)$ and $u_i = ter(S_i) \notin C$. We also assume that S_i is safe. Denote by B_i the set of inner points of the backward links of S_i, namely $B_i = \bigcup_{j \leq i} V(R_i^\circ)$.

Let $Z_{i+1} = \mathcal{Z}(u_i)$.

Assertion 4.16. *$u_i Z_{i+1}$ meets $V[\mathcal{Y}] \setminus B_i$.*

Assuming negation of the assertion, the alternating path obtained from S_i by adding to it the link $u_i Z_{i+1}$ shows that $ter(Z_{i+1}) \in SR$ (it is this argument for which we need to allow alternating paths to go through previous backward links). Then the path $u_i Z_{i+1}$ can serve as $T(u_i)$, to show that $u_i \in C$, a contradiction.

Let w_{i+1} be the first vertex on $u_i Z_{i+1}$ belonging to $V[\mathcal{Y}] \setminus B_1$. Let $Y_{i+1} = \mathcal{Y}(w_{i+1})$. Since $w_{i+1} \neq in(Z_{i+1})$ and $in[\mathcal{Y}] \subseteq in[\mathcal{Z}]$, we have $w_{i+1} \neq in(Y_{i+1})$. Let y_{i+1} be the vertex preceding w_{i+1} on Y_{i+1}.

Assertion 4.17. *$w_{i+1} \notin C$.*

Assuming for contradiction that $w_{i+1} \in C$, concatenating $T(y_{i+1})$ with $u_i Z_{i+1} w_{i+1}$ would yield a path $T(u_i)$, showing that $u_i \in C$. Note that $T(u_i)$ "ignores" the meeting with Y_i at w_i, but this is fine, since it needs not be $\mathcal{Y}$-committed.

Assertion 4.18. *$y_{i+1} \notin C$.*

Assuming for contradiction that $y_{i+1} \in C$, concatenating $T(y_{i+1})$ with the single edge link (y_{i+1}, w_{i+1}) and then with $u_i Z_{i+1} w_{i+1}$ would yield a path $T(u_i)$, showing that $u_i \in C$ (again, remember that $T(u_i)$ needs not be $\mathcal{Y}$-committed).

The last assertion, the fact that $w_{i+1} \notin B_i$, and the choice of u_i as the first vertex on Y_i not belonging to C, imply:

Assertion 4.19. *If $Y_{i+1} = Y_j = Y$ for some $j \leq i$ then $w_{i+1} >_Y w_j$.*

We continue the construction of the alternating paths S_i, adhering to the following two rules:

Rule 1: If $Y = Y_i \in \mathcal{Y}$ is met for the first time, we go on it backwards until we reach the first vertex u_i on Y not belonging to C.

Rule 2: If $Y = Y_i \in \mathcal{Y}$ has already been met, we go backwards on Y until we reach a vertex $w = w_j$ for some $j < i$, and let $u_i = w_j$.

Since by the induction hypothesis $w_j \notin C$ for $j < i$, by Assertion 4.17 when Rule 2 is applied we still have the condition $u_i \notin C$. Rule 2 guarantees that the alternating paths S_i constructed are safe. As noted above, the condition $ter(S_i) \notin C$ implies that the construction continues indefinitely, and generates an infinite $[\mathcal{Z}, \mathcal{Y}]$-alternating path S. In fact, S is safe, since Condition (1) of Definition 4.8 follows from the construction, while Condition (2) is true since the non-$\mathcal{Y}$ links in S come from $\mathcal{Z}$, which is f.c. This proves that possibility (1) of the lemma holds. □

Proof of Theorem 4.12 The connected components of the graph whose edge set is $E[\mathcal{Z}] \cup E[\mathcal{Y}]$ are countable. Hence we may assume that $\mathcal{Z}$ and $\mathcal{Y}$ are countable. Let $z_1, z_2, \ldots$ be an enumeration of $in[\mathcal{Z}] \setminus in[\mathcal{Y}]$. Applying Lemma 4.13 with $u = z_1$ we obtain a z_1-starting $[\mathcal{Z}, \mathcal{Y}]$-s.a.p Q_1, satisfying condition (1) or (2) of the lemma. If (1) is true, continue by applying the lemma to z_2. If (2) is true, denote the vertex v appearing in the lemma by v_1, and the (v, z_1)-$[\mathcal{Z}, \mathcal{Y}]$- alternating path by T_1. Then $\mathcal{Z}_1 = (\mathcal{Z} \triangle T_1)^{path}$ is a f.c. warp, with $in[\mathcal{Z}_1] = in[\mathcal{Z}] \setminus \{z_1\}$, $ter[\mathcal{Z}_1] = ter[\mathcal{Z}] \setminus \{v_1\}$.

(Recall that $(\mathcal{Z}\triangle T_1)^{path}$ is the warp obtained from $\mathcal{Z}\triangle T_1$ by removing its cycles. Such cycles might appear since T_1 is not required to be safe.) Apply now the lemma to the pair $(\mathcal{Z}_1, \mathcal{Y})$, with $u = z_2$.

Continuing this way, we obtain a sequence Q_i of z_i-starting $\mathcal{Y}$-s.a.p's, which are either infinite or end at distinct vertices of $ter[\mathcal{Z}]$, as promised in the theorem. □

Remark 4.20. The theorem applies also when $\mathcal{Z}$ is a fractured warp. Reducing the fractured case to the non-fractured case is done by duplicating those vertices which serve as both an initial point and a terminal point of paths from $\mathcal{Z}$, thus turning $\mathcal{Z}$ into a proper warp.

5. A Hall-type equivalent conjecture

In [3] Theorem 1.6 was shown to be equivalent to the following Hall-type conjecture:

Conjecture 5.1. *An unhindered web is linkable.*

Both implications in this equivalence are quite easy. To show how Theorem 1.6 implies Conjecture 5.1, suppose that Theorem 1.6 is true, and let $\mathcal{P}$ and S be as in the theorem. Then $\{Ps: \ P \in \mathcal{P}, s \in V(P) \cap S\}$ is a wave, and unless $\mathcal{P}$ is a linkage, it is also a hindrance. To prove the converse implication, take a $\preccurlyeq$-maximal wave $\mathcal{W}$ in Γ (see Lemma 3.20), and let $S = ter[\mathcal{E}(\mathcal{W})]$. By Lemma 3.26, Γ/S is loose, and in particular unhindered. Assuming that Conjecture 5.1 is true, the web Γ/S has therefore a linkage $\mathcal{L}$. Taking $\mathcal{P} = \mathcal{W} * \mathcal{L}$ then fulfils, together with S, the requirements of Theorem 1.6.

In fact, the above argument shows that the following is also equivalent to Theorem 1.6:

Conjecture 5.2. *A loose web is linkable.*

Here is a third equivalent formulation, generalizing Theorem 1.8:

Conjecture 5.3. *If Γ is unlinkable then there exists an A–B-separating set S which is linkable into A in $\overleftarrow{\Gamma}$, but A is not linkable into S in Γ.*

The main result of this paper is that Conjecture 5.1, and hence also Theorem 1.6, are true for general graphs. Let us thus re-state the conjecture, this time as a theorem:

Theorem 5.4. *An unhindered web is linkable.*

The rest of the paper is devoted to the proof of Theorem 5.4. The proof is divided into two stages. We first define a notion of a κ-*hindrance* for every regular uncountable cardinal κ, and show that the existence of a κ-hindrance implies the existence of a hindrance. Then we shall show that if a web is unlinkable then it contains either a hindrance or a κ-hindrance for some uncountable regular κ.

6. Safely linking one point

In this section we prove a result, whose key role was already mentioned in the introduction:

Theorem 6.1. *If Γ is unhindered then for every $a \in A$ there exists an a-B-path P such that $\Gamma - P$ is unhindered.*

Let us first outline the proof of the theorem in the case of countable graphs. This will serve two purposes: first, the main idea of the proof appears also in the general case; second, it will help to clarify the obstacle which arises in the uncountable case. A main ingredient in the proof is the following:

Lemma 6.2. *Let $Q \subseteq V \setminus (A \cup B)$, and let $\mathcal{U}$ be a wave in $\Gamma - Q$, such that*

$$N^+(Q) \setminus Q \subseteq RF_{\Gamma-Q}(\mathcal{U}) \ . \tag{2}$$

Then $\mathcal{U}$ is a wave in Γ.

Proof. Let P be an A–B-path. We have to show that P contains a vertex from $ter[\mathcal{U}]$. If P is disjoint from Q then, since $\mathcal{U}$ is a wave in $\Gamma - Q$, P contains a vertex from $ter[\mathcal{U}]$. If P meets Q then, since $Q \cap B = \emptyset$, there exists a vertex $y \in V(P) \cap N^+(Q) \setminus Q$. Choose y to be the last such vertex on P. By (2), the path yP then contains a vertex belonging to $ter[\mathcal{U}]$, as desired. □

Proof of Theorem 6.1 for countable webs. Enumerate all a-B-paths as $P_1, P_2, \ldots$. Assuming that the theorem fails, there exists a first vertex y_1 on P_1, such that $\Gamma - P_1y_1$ is hindered. Let $T_1 = P_1y_1 - y_1$. Then $\Gamma - T_1$ is unhindered. By Lemma 3.32, there exists a wave $\mathcal{W}_1$ in $\Gamma - T_1$ such that $y_1 \in ter[\mathcal{W}_1]$. Let i_2 be the first index (if such exists) such that P_{i_2} does not meet $V[\mathcal{W}_1]$. Let z be the last vertex on P_{i_2} lying on T_1, and let $P_2' = T_1zP_{i_2}$. We may assume $\Gamma - P_2'$ is unhindered and hence by Lemma 3.31, the web $\Gamma - T_1 - zP_{i_2}$ is also hindered, since it is obtained from $\Gamma - P_2'$ by removing finitely many vertices. Let y_2 be the first vertex on zP_{i_2} such that $\Gamma - T_1 - zP_{i_2}y_2$ is hindered, and let $T_2 = T_1 \cup (zP_{i_2}y_2 - y_2)$. By Lemma 3.32, there exists a wave $\mathcal{W}_2$ in $\Gamma - T_2$, such that $y_2 \in ter[\mathcal{W}_2]$.

Continuing this way, we obtain an ascending sequence of trees $(T_i : \ i < \mu)$ (where μ is either finite or ω), all rooted at a and directed away from a, and a sequence of waves $\mathcal{W}_i$ in $\Gamma - T_i$ disjoint from all trees T_j, such that every a-B-path contains a vertex separated by some $\mathcal{W}_i$ from B. Let $T = \bigcup_{i<\mu} T_i$ and $\mathcal{W} = \uparrow \mathcal{W}_i$. Each $\mathcal{W}_i$ is a wave in $\Gamma - T_i$ and hence also in $\Gamma - T$. Therefore Lemma 3.19 implies that $\mathcal{W}$ is a wave in $\Gamma - T$, separating from B at least one vertex from each a-B path. By Lemma 6.2, $\mathcal{W}$ is a wave in Γ, and since $a \notin in[\mathcal{W}]$, it is a hindrance, contradicting the assumption of the theorem, that Γ is unhindered. □

The difficulty in extending this proof beyond the countable case is that after ω steps the web $\Gamma - T_\omega$ may be hindered, and then we can not proceed with the same construction, since, for example, Lemma 3.32 is not applicable. Here is a brief outline of how this difficulty is overcome.

Why was the construction of the trees T_i necessary, and why wasn't it possible just to delete the initial parts of the paths P_i, and consider the waves (say) $\mathcal{U}_i$ resulting from those deletions? Because then each $\mathcal{U}_i$ lives in a different web, and it is impossible to combine the waves $\mathcal{U}_i$ to form one big wave. This we shall solve by taking quotient, instead of deleting vertices - as we saw in Lemma 3.29 it is then possible to combine the resulting waves. But then we obtain a wave which is not a wave in Γ, but in some quotient of it, namely it does not necessarily start in A, while for the final contradiction we need a wave (in fact, hindrance) in Γ itself. This we overcome by performing the proof in two stages. In the first we take quotients, and obtain a wave $\mathcal{W}$ "hanging in air" in Γ/X for some *countable* set X (keeping X countable is a key point in the proof). In the second stage we use the countability of X to delete its elements one by one, in a way similar to that used in the countable case, described above. This process will generate a wave $\mathcal{V}$, and the "arrow" concatenation of $\mathcal{V}$ and $\mathcal{W}$ will result in the desired wave in Γ.

Proof of Theorem 6.1. Construct inductively trees T_α rooted at a and directed away from a, as follows. The tree T_0 consists of the single vertex a. For limit ordinals β define $T_\beta = \bigcup_{\alpha<\beta} T_\alpha$. Assume that T_α is defined. Suppose first that there exists a vertex $x \in V \setminus (A \cup V(T_\alpha))$ such that $(u, x) \in E$ for some $u \in V(T_\alpha)$, and $\Gamma - a - F - x$ is unhindered for every finite subset F of $V(T_\alpha)$ not including a. In this case we choose such a vertex x, and construct $T_{\alpha+1}$ by adding x to $V(T_\alpha)$ and (u, x) to $E(T_\alpha)$. If no vertex x satisfying the above conditions exists, the process of definition is terminated at α, and we write $T = T_\alpha$.

The tree T thus constructed has the property that for every finite subset F of $V(T)$ not including a the web $\Gamma - a - F$ is unhindered, and T is maximal with respect to this property. Write $Y = N^+(V(T)) \setminus V(T)$. Then for every $y \in Y$ there exists a finite set $F_y \subseteq V(T) \setminus \{a\}$ such that $\Gamma - a - F_y - y$ is hindered. Thus, by Lemmas 3.32, there exists a wave $\mathcal{U}$ in $\Gamma - a - F_y$ with $y \in ter[\mathcal{U}]$. Let us now fix some maximal wave in $(\Gamma - a)/F_y$ and call it $\mathcal{A}_y$. Corollary 3.28 yields $y \in RF_\Gamma(\mathcal{A}_y)$.

Assuming that Theorem 6.1 fails, we have:

$$V(T) \cap B = \emptyset. \tag{3}$$

Call a vertex $t \in V(T)$ *bounded* if there exists a countable subset G_t of $V(T)$ and a wave $\mathcal{B} = \mathcal{B}_t$ in $(\Gamma - a)/G_t$ such that $t \in RF^\circ_\Gamma(\mathcal{B})$. Let Q be the set of non-bounded elements of $V(T)$. For every bounded vertex $t \in V(T)$ choose a fixed set G_t and a fixed wave $\mathcal{B}_t$ as above.

Let $\Gamma' = \Gamma - Q - a$. The core of the proof of Theorem 6.1 is in the following:

Proposition 6.3. *For every $y \in Y$ there exists a wave $\mathcal{U}_y$ in Γ' satisfying $y \in RF_{\Gamma'}(\mathcal{U}_y)$.*

Proof of the proposition: Let y be a fixed element of Y. We shall construct a countable subset X of $V(T) \setminus A$, and a wave $\mathcal{W}$ in $(\Gamma - a)/X$, having the following properties:

(a) $y \in RF(\mathcal{W})$.
(b) $F_z \subseteq X$ for every $z \in Y \cap V[\mathcal{W}\langle X\rangle]$.
(c) $G_t \subseteq X$ and $t \in RF^\circ_\Gamma(\mathcal{W})$ for every $t \in X \setminus Q$.
(d) $V[\mathcal{W}\langle X\rangle] \cap V(T) \subseteq X$.

The construction is by a "closing up" procedure. We construct an increasing sequence of sets X_i whose union is to be taken as X, and waves $\mathcal{W}_i$ in $(\Gamma - a)/X_i$ whose "$\uparrow$" limit will eventually be taken as $\mathcal{W}$, and at each step we take care of conditions (b) and (c), alternately, for all vertices $z \in Y \cap V[\mathcal{W}_i\langle X_i\rangle]$ and $t \in X_i \setminus Q$. We shall do this in steps, as follows.

We take $X_0 = F_y$ and let $\mathcal{W}_0 = \mathcal{A}_y$.

For every $i < \omega$ we then take $X_{i+1} = X_i \cup \bigcup_{z \in Y \cap V[\mathcal{W}_i\langle X_i\rangle]} F_z \cup \bigcup_{t \in X_i \setminus Q} G_t \cup (V[\mathcal{W}_i\langle X_i\rangle] \cap V(T))$

and let $\mathcal{W}_{i+1}$ be a maximal wave in $(\Gamma - a)/X_{i+1}$.

Let $X = \bigcup_{i<\omega} X_i$ and $\mathcal{W} = \uparrow_{i<\omega} (\mathcal{W}_i/X)$. Note that for every $z \in Y \cap V[\mathcal{W}\langle X\rangle]$ we have $z \in Y \cap V[\mathcal{W}_i\langle X_i\rangle]$ for some $i < \omega$. This implies $F_z \subseteq X_{i+1} \subseteq X$ proving condition (b). For every $t \in X \setminus Q$, we have $t \in X_i$ for some $i < \omega$ and hence $G_t \subseteq X_{i+1}$. Since $\mathcal{W}_{i+1}$ is a maximal wave in $(\Gamma - a)/X_{i+1}$, we have

$t \in RF^{\circ}_{(\Gamma-a)/X_{i+1}}(\mathcal{W}_{i+1})$. If $t \in \mathcal{E}(X)$, this implies $t \in RF^{\circ}_{(\Gamma-a)/X}(\mathcal{W}_{i+1}/X) \subseteq RF^{\circ}_{(\Gamma-a)/X}(\mathcal{W}) \subseteq RF^{\circ}_{\Gamma}(\mathcal{W})$, yeilding condition (c). Of course, if $t \in \mathcal{IE}(X)$, we still have $t \in RF^{\circ}_{\Gamma}(\mathcal{W})$. Conditions (d) is obviously taken care of by the construction. In view of Corollary 3.11, condition (a) has been taken care of by the fact that $\mathcal{W} \vec{\succcurlyeq} \mathcal{W}_1/X$.

By conditions (c) and (d), we have:

Assertion 6.4. *(i)* $ter[\mathcal{E}(\mathcal{W})\langle X\rangle] \cap V(T) \subseteq Q$.
(ii) $V[\mathcal{E}(\mathcal{W})\langle X\rangle] \cap Q \subseteq ter[\mathcal{E}(\mathcal{W})\langle X\rangle]$.

Proof. Let t be a vertex in $ter[\mathcal{E}(\mathcal{W})\langle X\rangle] \cap V(T)$. By condition (d) above, $t \in X$. Since by assumption $t \notin RF^{\circ}(\mathcal{W})$, by condition (c) it follows that $t \in Q$. This proves (i).

To prove part (ii), assume that $q \in (Q \cap V[\mathcal{W}\langle X\rangle]) \setminus ter[\mathcal{E}(\mathcal{W})\langle X\rangle]$. By the self roofing lemma (Lemma 3.7), it follows that $q \in RF^{\circ}(\mathcal{W})$. But, since $\mathcal{W}$ is a wave in Γ/X, and X is countable, this contradicts the fact that $q \in Q$. □

Let $\mathcal{W}'$ be obtained from $\mathcal{E}(\mathcal{W})$ by the removal of all paths ending at Q. By Assertion 6.4 (ii), $\mathcal{W}'$ is a wave in $(\Gamma/X) - Q - a$, and by condition (a) we have $y \in RF_{\Gamma'}(\mathcal{W}')$. Thus $\mathcal{W}'$ has almost all properties required from the wave $\mathcal{U}$ in the proposition, the only problem being that we are looking for a wave $\mathcal{U}$ in $\Gamma - Q - a$, not in $\Gamma/X - Q - a$. We now wish to "bring $\mathcal{W}'$ to the ground", namely make it start at A, not at $A \cup X$.

To achieve this goal, we enumerate the vertices of X as $x_1, x_2, \ldots$, and start deleting them one by one - this time, real deletion, not the quotient operation. Let $k_1 = 1$, delete $x_{k_1} = x_1$, and choose a maximal wave $\mathcal{V}_1$ in $\Gamma - a - x_1$. Next choose the first vertex x_{k_2} not belonging to $V[\mathcal{V}_1]$ (if such exists), take a maximal wave $\mathcal{V}'_2$ in $\Gamma - a - \{x_{k_1}, x_{k_2}\}$, and define $\mathcal{V}_2 = \mathcal{V}_1 \frown \mathcal{V}'_2$. Then choose the first k_3 such that $x_{k_3} \notin V[\mathcal{V}_2]$ (if such exists), take a maximal wave $\mathcal{V}'_3$ in $\Gamma - a - \{x_{k_1}, x_{k_2}, x_{k_3}\}$, and define $\mathcal{V}_3 = \mathcal{V}_2 \frown \mathcal{V}'_3$. If the process terminates after m steps for some finite m, let $\mathcal{V} = \mathcal{V}_m$. Otherwise, let $\mathcal{V} = \uparrow_{k<\omega} \mathcal{V}_k$. Let $\theta = \omega$ if this process lasts ω steps, and $\theta = m+1$ if it terminates after m steps for some finite number m. For $i < \theta$ denote the set $\{x_{k_1}, x_{k_2}, \ldots, x_{k_i}\}$ by R_i, and write $R = \{x_{k_1}, x_{k_2}, x_{k_3} \ldots\}$.

Our goal now is to show that $\mathcal{V} \frown \mathcal{W}'$ is a wave in Γ'. This will be done by applying Lemma 3.15 with Γ replaced by Γ', the wave $\mathcal{U}$ replaced by $\mathcal{V}$, the wave $\mathcal{W}$ replaced by $\mathcal{W}'$, the set X in the lemma replaced by R and the set Z replaced by X. We already know that $\mathcal{W}'$ is a wave in Γ'/X. We need to show that $\mathcal{V}$ is a wave in $\Gamma' - R$ and every path in $\mathcal{W}'$ meets $RF_{\Gamma'-R}(\mathcal{V})$. Note that we already know that $\mathcal{V}$ is a wave in $\Gamma - R$. Therefore, in order to show that it is a wave in $\Gamma' - R$, we only need to prove it does not meet Q. Also note that following Remark 3.16, in order to show that every path in $\mathcal{W}'$ meets $RF_{\Gamma'-R}(\mathcal{V})$, it is enough to consider only paths starting at X.

Recall that $\mathcal{V}'_i$ is a $\leq$-maximal wave in $\Gamma - a - R_i$ and by Lemma 3.17, so is $\mathcal{V}_i$.

We also have

Assertion 6.5. $V(T) \cap ter[\mathcal{V}] = \emptyset$.

Proof. If $t \in V(T) \cap ter[\mathcal{V}]$ then $t = ter(P)$ for some $P \in \mathcal{V}_i$ for some i. But then, the wave $\mathcal{V}_i \setminus \{P\}$ is a hindrance in $\Gamma - \{a, x_{k_1}, x_{k_2}, \ldots, x_{k_i}, t\}$, contradicting the fact that the deletion of any finite subset of $V(T)$ does not generate a hindrance. □

Assertion 6.6. $V[\mathcal{V}] \cap Q = \emptyset$.

Proof. Suppose, for contradiction, that $V[\mathcal{V}] \cap Q \neq \emptyset$. Then there exists $i < \theta$ and $q \in Q$ such that $q \in V[\mathcal{V}_i]$. By Assertion 6.5, $q \notin ter[\mathcal{V}_i]$, and since $\mathcal{V}_i$ is a wave in $\Gamma - a - R_i$, by the self roofing lemma (Lemma 3.7) $q \in RF^{\circ}_{\Gamma-a-R_i}(\mathcal{V}_i)$. By Lemma 3.28 it follows that $q \in RF^{\circ}_{\Gamma}(\mathcal{U})$, where $\mathcal{U}$ is a maximal wave in $(\Gamma - a)/R_i$. But this contradicts the definition of Q. □

Remark 6.7. As pointed out by R. Diestel, Assertion 6.6 is not essential for the argument that follows, since by the definition of Q we have: $V[\mathcal{V}] \cap Q \subseteq ter[\mathcal{V}]$. Thus we could replace $\mathcal{V}$ by $\mathcal{V}' = \mathcal{V} \setminus \mathcal{V}\langle Q\rangle$, and the argument below would remain valid. But since in fact $\mathcal{V}' = \mathcal{V}$, we chose the longer, but more informative, route.

By Assertion 6.6 $\mathcal{V}$ is a wave in $\Gamma' - R$.

Assertion 6.8. *If $z \in Y \cap V[\mathcal{W}\langle X\rangle]$ then $z \in RF_{\Gamma'-R}(\mathcal{V})$.*

Proof. By (b) we have $F_z \subseteq X$. Let $n < \omega$ be chosen so that $X' = \{x_1, \ldots, x_n\} \supseteq F_z$. Since $\Gamma - a - X'$ is unhindered and $\Gamma - a - X' - z$ is hindered, by Lemma 3.32 there exits a wave $\mathcal{Z}$ in $\Gamma - a - X'$ with $z \in ter[\mathcal{Z}]$. Let i be maximal such that $R_i \subseteq X'$. By the maximality property of i we have $X' \setminus R_i \subseteq V[\mathcal{V}_i] \subseteq RF_{\Gamma-a-R_i}(\mathcal{V}_i)$.

We now note that $\mathcal{V}_i$ is a wave in $\Gamma - a - R_i$ and $\mathcal{Z}$ is a wave in $\Gamma - a - X'$. Hence we can conclude that $\mathcal{V}_i \frown \mathcal{Z}$ is a wave in $\Gamma - a - R_i$, by applying Lemma 3.14 (with Γ replaced by $\Gamma - a - R_i$, the wave $\mathcal{U}$ replaced by $\mathcal{V}_i$, the wave $\mathcal{W}$ replaced by $\mathcal{Z}$ and Y replaced by $X' \setminus R_i$). By the maximality of $\mathcal{V}_i$ we have $\mathcal{V}_i = \mathcal{V}_i \frown \mathcal{Z}$. This implies that $z \in RF_{\Gamma-a-R_i}(\mathcal{V}_i)$. Since $\mathcal{V} \succcurlyeq \mathcal{V}_i$ and $R \cup Q \supseteq R_i$ we have $z \in RF_{\Gamma-a-Q-R}(\mathcal{V})$. □

Corollary 6.9. *Every path in $\mathcal{W}'$ meets $RF_{\Gamma'-R}(\mathcal{V})$.*

Proof. Let W be a path in $\mathcal{W}'$ and let $w = in(W)$. Then either $w \in A$ or $w \in X$. If $w \in A$ then since $\mathcal{V}$ is a wave in $\Gamma' - R$, we have $w \in RF_{\Gamma'-R}(\mathcal{V})$. If $w \in X$ then let $t = ter(W)$. Since $\mathcal{W}'$ was obtained from $\mathcal{W}$ by removing paths ending at Q, we have $t \notin Q$. By Assertion 6.4(i), we now have $t \notin V(T)$. Let z be the first vertex in W outside $V(T)$. Then $z \in Y$ and by assertion 6.8 we have $z \in F_{\Gamma'-R}(\mathcal{V})$. □

Define: $\mathcal{U}_y = \mathcal{V} \frown \mathcal{W}'$. Apply Lemma 3.15 with Γ replaced by Γ', the wave $\mathcal{U}$ replaced by $\mathcal{V}$, the wave $\mathcal{W}$ replaced by $\mathcal{W}'$, the set X in the lemma replaced by R and the set Z replaced by X. Corollary 6.9 asserts that indeed every path in $\mathcal{W}'$ meets $RF_{\Gamma'-R}(\mathcal{V})$ as needed to apply the lemma. The lemma yields that the warp $\mathcal{U}_y$ is a wave in Γ'. This completes the proof of Proposition 6.3.

To end the proof of Theorem 6.1, let $\mathcal{U} = \uparrow_{y \in Y} \mathcal{U}_y$. Then $\mathcal{U}$ separates Y from B. By Lemma 6.2 it follows that $\mathcal{U}$ is a wave in Γ, and since it does not contain a as an initial vertex of a path, it is a hindrance in Γ. This contradicts the assumption that Γ is unhindered. □

7. κ-LADDERS AND κ-HINDRANCES

7.1. **Stationary sets.** As is customary in set theory, an ordinal is taken as the set of ordinals smaller than itself, and a cardinal κ is identified with the smallest ordinal of cardinality κ. An uncountable cardinal λ is called *singular* if there exists a sequence $(\nu_\alpha : \alpha < \mu)$ of ordinals, whose limit is λ, where all ν_α, as well as μ, are smaller than λ. The smallest singular cardinal is $\aleph_\omega$, which is the limit of $(\aleph_i : i < \omega)$. A singular cardinal is necessarily a limit cardinal, namely it must be of the form $\aleph_\theta$ for some limit ordinal θ. On the other hand, ZFC (assuming its consistency) has models in which there exist non-singular limit cardinals.

A non-singular cardinal is called *regular*.

The main set-theoretic notion we shall use is that of *stationary sets*. A subset of an uncountable regular cardinal κ is called *unbounded* if its supremum is κ, and *closed* if it contains the supremum of each of its bounded subsets. A subset of κ is called *stationary* (or κ-stationary) if it intersects every closed unbounded subset of κ. A function f from a set of ordinals to the ordinals is called *regressive* if $f(\alpha) < \alpha$ for all α in the domain of f. A basic fact about stationary sets is Fodor's lemma:

Theorem 7.1. *If κ is regular and uncountable, Φ is a κ-stationary set, and $f : \Phi \to \kappa$ is regressive, then there exist a stationary subset Φ' of Φ and an ordinal β such that $f(\phi) = \beta$ for all $\phi \in \Phi'$.*

Fodor's lemma implies that stationary sets are in some sense "big". This is expressed also in the following:

Lemma 7.2. *If Ξ_α, $\alpha < \lambda$ are non-stationary, and $\lambda < \kappa$, then $\bigcup_{\alpha<\lambda} \Xi_\alpha$ is non-stationary.*

This is another way of saying that the intersection of fewer than κ closed unbounded sets is closed and unbounded.

7.2. **κ-ladders.** The tool used in the proof of Theorem 5.4 in the uncountable case is *κ-ladders*, for uncountable regular cardinals κ. A κ-ladder $\mathcal{L}$ is a sequence of "rungs" $(R_\alpha : \alpha < \kappa)$. At each step α we are assuming that a warp $\mathcal{Y}_\alpha = \mathcal{Y}_\alpha(\mathcal{L})$ in Γ is defined, by the previous rungs of $\mathcal{L}$. For each $\alpha \geq 0$, assuming $\mathcal{Y}_\alpha$ is defined, we let $\Gamma_\alpha = \mathcal{E}(\Gamma/\mathcal{Y}_\alpha)$.

The warp $\mathcal{Y}_0$ is defined as $\langle A \rangle$, and for limit ordinals α, we let $\mathcal{Y}_\alpha = \uparrow_{\theta<\alpha} \mathcal{Y}_\theta$.

For successor ordinals $\alpha + 1$, the warp $\mathcal{Y}_{\alpha+1}$ is defined by $\mathcal{Y}_\alpha$ and by the rung R_α, the latter being chosen as follows. A first constituent of R_α is a (possibly trivial) wave $\mathcal{W}_\alpha$ in Γ_α. If the set $V(\Gamma_\alpha) \setminus (A(\Gamma_\alpha) \cup V[\mathcal{W}_\alpha])$ is non-empty, then R_α consists also of a vertex y_α in this set. The warp $\mathcal{Y}_{\alpha+1}$ is defined in this case as $\mathcal{Y}_\alpha \frown \mathcal{W}_\alpha \cup \langle y_\alpha \rangle$. If $V(\Gamma_\alpha) \setminus (A(\Gamma_\alpha) \cup V[\mathcal{W}_\alpha]) = \emptyset$, then $\mathcal{Y}_{\alpha+1}$ is defined as $\mathcal{Y}_\alpha \frown \mathcal{W}_\alpha$. In this case all consecutive rungs will consist just of the trivial wave, meaning that the ladder will "mark time", without changing.

We also wish to keep track of the steps in which a new hindrance emerges in the ladder. This is done by keeping record of subsets $\mathcal{H}_\alpha$ of $\mathcal{Y}_\alpha$. These sets are not uniquely defined by $\mathcal{L}$, but to simplify notation we assume that the ladder comes with a fixed choice of such sets, which is subject to the following conditions.

We define $\mathcal{H}_0 = \emptyset$. If $\mathcal{IE}(\mathcal{Y}_{\alpha+1}) \setminus \mathcal{H}_\alpha \neq \emptyset$ we pick a (possibly unending) path H in this set, write $H_\alpha = H$, and $\mathcal{H}_{\alpha+1} = \bigcup_{\theta<\alpha} \mathcal{H}_\theta \cup \{H\}$.

If $\mathcal{IE}(\mathcal{Y}_{\alpha+1}) \setminus \mathcal{H}_\alpha \neq \emptyset$ we let $\mathcal{H}_{\alpha+1} = \mathcal{H}_\alpha$. For limit α we define $\mathcal{H}_\alpha = \bigcup_{\theta<\alpha} \mathcal{H}_\theta$.

Remark 7.3. Note that it is possible that $\bigcup_{\alpha<\kappa} \mathcal{H}_\alpha \neq \mathcal{IE}(\mathcal{Y})$, namely that we never exhaust all of $\mathcal{IE}(\mathcal{Y})$.

Since a path in $\mathcal{H}_\alpha$ is inessential in $\mathcal{Y}_\alpha$, it will never "grow" in any later stage β, and hence we have:

Lemma 7.4. *$\mathcal{H}_\alpha \subseteq \mathcal{IE}(\mathcal{Y}_\beta)$ for all $\beta \geq \alpha$.*

The set of ordinals α for which $\mathcal{IE}(\mathcal{Y}_{\alpha+1}) \setminus \mathcal{H}_\alpha \neq \emptyset$ is denoted by $\Phi(\mathcal{L})$. As noted, $\Phi(\mathcal{L})$ is not uniquely defined by $\mathcal{L}$ itself, and is dependent on the choice of the sets $\mathcal{H}_\alpha$.

Example 7.5. Let $|A| = \aleph_0,\ B = \emptyset,\ V(\Gamma) = A$, and choose $\kappa = \aleph_1$. Since Γ_1 is defined as $\mathcal{E}(\Gamma/\langle A\rangle)$, it is empty (i.e., Γ_1 has no vertices), and $\mathcal{Y}_\alpha = \mathcal{IE}(\mathcal{Y}_\alpha) = \langle A\rangle$ for all $1 \leq \alpha < \aleph_1$. The paths $(a),\ a \in A$ can be chosen as H_α in any order, and thus $\Phi(\mathcal{L})$ can be any countable ordinal.

We write $\Phi^\infty(\mathcal{L})$ for the set of those $\alpha \in \Phi(\mathcal{L})$ for which $\mathcal{IE}(\mathcal{Y}_{\alpha+1}) \setminus \mathcal{H}_\alpha$ contains an unending path, and Φ^{fin} for $\Phi(\mathcal{L}) \setminus \Phi^\infty(\mathcal{L})$.

Let $\Phi_h(\mathcal{L}) = \{\alpha \mid \mathcal{W}_\alpha \text{ is a hindrance}\}$, and $\Phi_h^\infty(\mathcal{L}) = \{\alpha \mid \mathcal{Y}_\alpha^\infty \setminus \bigcup_{\theta<\alpha} \mathcal{Y}_\theta^\infty \neq \emptyset\}$. Unlike $\Phi(\mathcal{L})$, the set $\Phi_h(\mathcal{L})$ is determined by $\mathcal{L}$. The difference between the two sets is that the ordinals in $\Phi_h(\mathcal{L})$ are "newly hindered", namely there is a hindered vertex generated at that stage, whereas the fact that $\alpha \in \Phi(\mathcal{L})$ means that not all hindered vertices generated so far have been "taken into account", in the sense of being included in $\mathcal{H}_\alpha$. In Example 7.5 $\Phi_h(\mathcal{L}) = \{0\}$.

Lemma 7.6. $\Phi_h(\mathcal{L}) \subseteq \Phi(\mathcal{L})$.

Proof. Suppose that $\alpha \in \Phi_h(\mathcal{L})$. We shall show that $\mathcal{IE}(\mathcal{Y}_{\alpha+1}) \setminus \mathcal{H}_\alpha \neq \emptyset$, which will imply the desired inclusion result. Let x be a vertex in $A(\Gamma_\alpha) \setminus in[\mathcal{W}_\alpha]$. Then $x = ter(P)$ for some $P \in \mathcal{E}(\mathcal{Y}_\alpha)$. By the definition of $\mathcal{H}_\alpha$, we have $P \notin \mathcal{H}_\alpha$. By the definition of a wave, $ter[\mathcal{W}_\alpha]$ is separating in Γ_α and thus also in Γ. The set $ter[\mathcal{Y}_\alpha \frown \mathcal{W}_\alpha] \setminus \{x\}$ contains $ter[\mathcal{W}_\alpha]$ and is hence separating as well. Therefore $P \in \mathcal{IE}(\mathcal{Y}_\alpha \frown \mathcal{W}_\alpha)$. Thus $\mathcal{IE}(\mathcal{Y}_\alpha \frown \mathcal{W}_\alpha) \setminus \mathcal{H}_\alpha \neq \emptyset$, meaning that R_α is hindered. □

Lemma 7.7. $\Phi_h^\infty(\mathcal{L}) \subseteq \Phi^\infty(\mathcal{L})$.

Proof. Let α be an ordinal in $\Phi_h^\infty(\mathcal{L})$, and let P be a path witnessing this, namely $P \in \mathcal{Y}_\alpha^\infty \setminus \bigcup_{\theta<\alpha} \mathcal{Y}_\theta^\infty$. Then $P \notin \bigcup_{\theta<\alpha} \mathcal{Y}_\theta$, and since $\mathcal{H}_\alpha \subseteq \bigcup_{\theta<\alpha} \mathcal{IE}(\mathcal{Y}_\theta)$, this implies that $P \in \mathcal{IE}(\mathcal{Y}_\alpha) \setminus \mathcal{H}_\alpha$. □

The following is obvious from the way the sets $\mathcal{H}_\alpha$ are chosen:

Lemma 7.8. *If* $|\mathcal{IE}(\mathcal{Y}_\alpha)| \geq \kappa$ *for some* $\alpha < \kappa$, *then* $\Phi(\mathcal{L}) \supseteq [\alpha, \kappa)$.

Notation 7.9. Write $T_\alpha = T_\alpha(\mathcal{L})$ for $A(\Gamma_\alpha)$. The warp $\mathcal{Y}_\kappa$ is denoted by $\mathcal{Y} = \mathcal{Y}(\mathcal{L})$. For $\alpha \in \Phi^{fin}(\mathcal{L})$ denote $ter(H_\alpha)$ by x_α. The set $\{y_\alpha :\ \alpha < \kappa\}$ is denoted by $Y(\mathcal{L})$, and for every $\beta \leq \kappa$ write $Y_\beta(\mathcal{L})$ for $\{y_\alpha :\ \alpha < \beta\}$. The set $\{x_\alpha \mid \alpha \in \Phi^{fin}(\mathcal{L})\}$ is denoted by $X^{fin}(\mathcal{L})$.

The definitions clearly imply:

Lemma 7.10. T_α *is* A–B*-separating for all* $\alpha < \kappa$. *If* $\alpha < \beta$ *then* $T_\alpha \subseteq RF(T_\beta)$.

By the definition of Γ_α as $\mathcal{E}(\Gamma/\mathcal{Y}_\alpha)$ we have:

Lemma 7.11. T_α *is a minimal* A–B*-separating for all* $\alpha < \kappa$.

Define $RF(\mathcal{L}) = \bigcup_{\theta<\kappa} RF(T_\theta)$ and $RF^\circ(\mathcal{L}) = \bigcup_{\theta<\kappa} RF^\circ(T_\theta)$.

Also write $\Gamma^\alpha = \Gamma[RF(T_\alpha)]$, which means that $D(\Gamma^\alpha)$ (the digraph of Γ^α) is $\Gamma[RF(T_\alpha)]$, $A(\Gamma^\alpha) = A$ and $B(\Gamma^\alpha) = T_\alpha$.

For $\alpha < \beta$ let Γ_α^β be the part of Γ between T_α and T_β, namely $V(\Gamma_\alpha^\beta) = V(\Gamma_\alpha[RF_{\Gamma_\alpha}(T_\beta)])$, $D(\Gamma_\alpha^\beta) = D(\Gamma_\alpha[RF_{\Gamma_\alpha}(T_\beta)])$, $A(\Gamma_\alpha^\beta) = T_\alpha,\ B(\Gamma_\alpha^\beta) = T_\beta$.

Notation 7.12. We shall write $V^\alpha = V^\alpha(\mathcal{L})$ for $V(\Gamma^\alpha)$, and V_α for $V(\Gamma_\alpha)$, namely $V^\alpha = RF(T_\alpha)$ and $V_\alpha = V(\Gamma) \setminus RF^\circ(T_\alpha)$.

Notation 7.13. Let $\Phi_G(\mathcal{L}) = \{\alpha \in \Phi(\mathcal{L}) \mid in(H_\alpha) \in A\}$ and $\Phi_H(\mathcal{L}) = \Phi(\mathcal{L}) \setminus \Phi_G(\mathcal{L})$ (The "G" stands for "grounded" and the "H" stands for "hanging in air").

Throughout the proof we shall construct again and again ladders, which will all be denoted by $\mathcal{L}$. In all these cases we shall use the following:

Convention 7.14. We shall denote $\mathcal{Y}(\mathcal{L})$, for the ladder $\mathcal{L}$ considered at that point, by $\mathcal{Y}$. We shall also write T_α for $T_\alpha(\mathcal{L})$, Y_α for $Y_\alpha(\mathcal{L})$, Φ for $\Phi(\mathcal{L})$, and so on.

Lemma 7.15. $\Phi_H(\mathcal{L})$ *is non-stationary.*

Proof. For $\alpha \in \Phi_H(\mathcal{L})$ we have $in(H_\alpha) = y_\beta$ for some $\beta < \alpha$. The function $f(\alpha) = \beta$ defined in this way is a regressive injection from $\Phi_H(\mathcal{L})$ to κ. Thus, by Fodor's lemma, $\Phi_H(\mathcal{L})$ is not stationary. □

The following is obvious:

Lemma 7.16. *A vertex $v \in V$ belongs to $RF(\mathcal{L}) \setminus RF^\circ(\mathcal{L})$ if and only if there exists $\beta < \kappa$ such that $v \in T_\alpha$ for all $\alpha \geq \beta$.*

Lemma 7.17. *Let Q be a $\mathcal{Y}$-alternating path, and assume that $in(Q) \in RF^\circ(T_\alpha)$. Then:*

1. *$V(Q) \subseteq RF(T_\alpha)$, and:*
2. *If $in(Q) = x_\alpha$ and $ter(Q) = y_\beta$, then $\beta < \alpha$.*

Proof. Write $z = in(Q)$. Using the same notation as in Definition 4.2, write Q as $(z = z_0, F_1, u_1, R_1, z_1, F_2, u_2, R_2, z_2 ...)$, where F_i are forward paths, namely using edges not belonging to $E[\mathcal{Y}]$, R_i are backward paths, namely using edges of $E[\mathcal{Y}]$, u_i are vertices on paths from $\mathcal{Y}$ at which Q switches from forward to backward direction, and z_i vertices at which Q switches from backward to forward direction. Since $z \in RF^\circ(T_\alpha)$, and T_α separates $V[\mathcal{L}]$ from B, F_1 is contained in $RF(T_\alpha)$. Possibly $u_1 \in T_\alpha$, but since R_1 goes backwards, $z_1 \in RF^\circ(T_\alpha)$. Thus F_2 is contained in $RF(T_\alpha)$. By an inductive argument following these steps we obtain part 1 of the lemma.

If $ter(Q) = y_\beta$, then by part (1), $y_\beta \in RF(T_\alpha)$. But $y_\beta \in V(\Gamma_\beta) \setminus A(\Gamma_\beta) = V(\Gamma) \setminus RF(T_\beta)$. Therefore $RF(T_\alpha) \setminus RF(T_\beta) \neq \emptyset$, and hence $\beta < \alpha$. □

Write $\zeta(\alpha)$ for the minimal ordinal at which H_α emerges as an inessential path, namely the minimal ordinal β such that $H_\alpha \in \mathcal{IE}(\mathcal{Y}_\beta)$. The choice of H_α implies:

Lemma 7.18. $\zeta(\alpha) \leq \alpha$ *for all* $\alpha \in \Phi(\mathcal{L})$.

Since $H_\alpha \in \mathcal{IE}(\mathcal{Y}_{\zeta(\alpha)})$, we have:

Lemma 7.19. $x_\alpha \in RF^\circ(T_{\zeta(\alpha)})$ *for every* $\alpha \in \Phi^{fin}(\mathcal{L})$.

Combined with Lemma 7.18, this yields:

Lemma 7.20. $x_\alpha \in RF^\circ(T_\alpha)$ *for every* $\alpha \in \Phi^{fin}(\mathcal{L})$.

7.3. **κ-hindrances.** Ordinals in $\Phi(\mathcal{L})$ are "troublesome", witnessing as they do the existence of hindrances. Thus, if $\Phi(\mathcal{L})$ is "large" then the ladder may pose a problem for linkability of Γ. And now we know what "large" should be: stationary. This is the origin of the following definition:

Definition 7.21. If $\Phi(\mathcal{L})$ is κ-stationary, then $\mathcal{L}$ is called a *κ-hindrance.*

Lemmas 7.15 and 7.2 yield together:

Lemma 7.22. *If $\mathcal{L}$ is a κ-hindrance then $\Phi_G(\mathcal{L})$ is stationary.*

Example 7.23. Let A be a set of size $\aleph_1$, B a set of size $\aleph_0$, let D be the complete directed graph on (A, B), namely $E(D) = A \times B$, and let $\Gamma = (D, A, B)$. We define an $\aleph_1$ ladder in Γ, as follows. Order B as $(b_\alpha \mid \alpha < \omega)$ and A as $(a_\alpha \mid \alpha < \omega_1)$.

For $\alpha < \omega$ let $\mathcal{W}_\alpha$ be the trivial wave, and $y_\alpha = b_\alpha$. Then for all such α we have $\Gamma_\alpha = \Gamma/\{b_i \mid i < \alpha\}$ and $\mathcal{H}_\alpha = \emptyset$. At the ω step we have $\mathcal{Y}_\omega = \langle A \cup B\rangle$, $\Gamma_\omega = \Gamma/B = ((B, \emptyset), B, B)$ and $\mathcal{H}_\omega = \emptyset$. Note that all the singleton paths in $\langle A\rangle$ are inessential in $\mathcal{Y}_\omega$.

For $0 \leq \alpha \leq \aleph_1$ let $R_{\omega+\alpha}$ consist of the inessential singleton path $H_{\omega+\alpha} = (a_\alpha)$. We then have $\mathcal{Y}_{\omega+\alpha} = \langle A \cup B\rangle$, $\Gamma_{\omega+\alpha} = ((B, \emptyset), B, B)$ and $\mathcal{H}_{\omega+\alpha} = \langle\{a_\theta \mid \theta < \alpha\}\rangle$.

Thus $\Phi(\mathcal{L}) = [\omega, \aleph_1)$, which is stationary, and hence $\mathcal{L}$ is an $\aleph_1$-hindrance.

Example 7.24 (accommodated from [11]). Let κ be an uncountable regular cardinal, and Ψ a κ-stationary set. Let $A = \{a_\alpha \mid \alpha \in \Psi\}$, $B = \{b_\alpha \mid \alpha < \kappa\}$, and let D be the directed graph whose vertex set is $A \cup B$ and whose edge set is $E = \{(a_\alpha, b_\beta) \mid \beta < \alpha\}$. Let $\Gamma = (D, A, B)$.

By Fodor's lemma, Γ is unlinkable.

Define a κ-ladder in Γ as follows. For all $\alpha < \kappa$ let $y_\alpha = b_\alpha$ and let $\mathcal{W}_\alpha$ be the trivial wave. Define sets $\mathcal{H}_\alpha$ by adding to $\mathcal{H}_\alpha$, for each $\alpha \in \Psi$, the singleton inessential path $H_\alpha = (a_\alpha)$. Here we have $\mathcal{Y}_\alpha = \langle A \cup \{b_\theta \mid \theta < \alpha\}\rangle$ and the path (a_β) is inessential in it for every $\beta \leq \alpha$. Since Ψ is stationary, this is a κ-hindrance.

Example 7.25. The following example shows the role of infinite paths in κ-hindrances. Let Ψ be an $\aleph_1$-stationary set all of whose element are limit ordinals (e.g., Ψ can be the set of *all* countable limit ordinals). For every $\alpha \in \Psi$, let $(\eta_i^\alpha \mid i < \omega)$ be an ascending sequence converging to α, where $\eta_0^\alpha = 0$.

Let $C = \{c_i^\alpha \mid \alpha \in \Psi,\ i < \omega\}$, $B = \{b_\alpha : \alpha < \omega_1\}$, let A be the subset of C $A = \{c_0^\alpha \mid \alpha \in \Psi\}$, let D be the directed graph whose vertices are $C \cup B$ and whose edges are $E = \{(c_i^\alpha, c_{i+1}^\alpha) \mid \alpha \in \Psi,\ i < \omega\} \cup \{(c_i^\alpha, c_j^\beta) \mid \alpha, \beta \in \Psi,\ i, j < \omega,\ \beta < \alpha,\ \eta_i^\alpha \geq \eta_j^\beta\} \cup \{(c_i^\alpha, b_\beta) \mid \alpha \in \Psi,\ i < \omega,\ \beta \leq \eta_i^\alpha\}$ and and let $\Gamma = (D, A, B)$.

Again, by Fodor's lemma, Γ is unlinkable.

We can construct an $\aleph_1$-ladder $\mathcal{L}$ on Γ by taking $y_\alpha = b_\alpha$ and $\mathcal{W}_\alpha = \{(b_\beta) \mid \beta < \alpha\} \cup \{(c_i^\beta, c_{i+1}^\beta) \mid \eta_{i+1}^\beta = \alpha\} \cup \{(c_i^\beta) \mid \eta_i^\beta < \alpha < \eta_{i+1}^\beta\}$. For $\alpha \in \Psi$, the concatenation of these waves forms an infinite path $(c_0^\alpha, c_1^\alpha, c_2^\alpha, c_3^\alpha, \ldots)$ in $\mathcal{Y}_\alpha$. We can take this path as H_α.

This yields $\Phi(\mathcal{L}) = \Psi$ and therefore $\mathcal{L}$ is an $\aleph_1$-hindrance.

Lemma 7.26. *If Γ does not contain a κ-hindrance then for every κ-ladder $\mathcal{L}$ and every $\alpha < \kappa$ there holds $|\mathcal{Y}_\alpha\langle\sim T_\alpha\rangle| < \kappa$.*

Proof. A path $P \in \mathcal{Y}_\alpha$ not meeting T_α belongs to $\mathcal{IE}(\mathcal{Y}_\alpha)$. Hence, if $|\mathcal{Y}_\alpha\langle\sim T_\alpha\rangle| \geq \kappa$ then $|\mathcal{IE}(\mathcal{Y}_\alpha)| \geq \kappa$, and hence by Lemma 7.8 $\mathcal{L}$ is a κ-hindrance. $\square$

The following lemma is not essential for the discussion to follow, but its understanding may clarify the nature of κ-hindrances. It says that Lemmas 7.6, 7.7 and 7.8 summarize all reasons for $\mathcal{L}$ to be a κ-hindrance:

Lemma 7.27. *A κ-ladder $\mathcal{L}$ is a κ-hindrance if and only if either:*
(i) $\Phi_h(\mathcal{L}) \cup \Phi_h^{\infty}(\mathcal{L})$ is stationary, or:
(ii) $|\mathcal{IE}(\mathcal{Y}_\alpha)| \geq \kappa$ for some $\alpha < \kappa$.

This means, among other things, that although $\Phi(\mathcal{L})$ is not uniquely determined by $\mathcal{L}$, whether it is stationary or not is determined by $\mathcal{L}$ alone. Namely, $\mathcal{L}$ being a κ-hindrance is independent of the order by which the paths H_α are chosen. The lemma also clarifies why we need to work with $\Phi(\mathcal{L})$ rather than $\Phi_h(\mathcal{L})$: because of the possible occurrence of case (ii).

Proof of Lemma 7.27: In view of Lemmas 7.6, 7.7 and 7.8, it remains to be shown that if $\Phi(\mathcal{L})$ is stationary, then one of conditions (i) and (ii) is true. By Lemma 7.18 $\zeta(\alpha) \leq \alpha$ for all α. If the set $\{\alpha \mid \zeta(\alpha) = \alpha\}$ is stationary, then (i) holds. Otherwise, assuming $\Phi(\mathcal{L})$ is stationary, by Fodor's lemma there exist a stationary subset $\Phi' \subseteq \Phi(\mathcal{L})$ and an ordinal $\beta < \kappa$, such that $\zeta(\alpha) = \beta$ for every $\alpha \in \Phi'$. By the definition of ζ this implies that $|\mathcal{IE}(\mathcal{Y}_\beta)| \geq \kappa$, proving (ii). □

Lemma 7.28. *Let $\mathcal{L}$ be a κ-ladder that is not a κ-hindrance, and let Σ be a closed unbounded set avoiding $\Phi(\mathcal{L})$. Then for every $P \in \mathcal{Y}(\mathcal{L})$ the set $\Sigma(P) = \{\alpha \in \Sigma \mid T_\alpha \cap V(P) \neq \emptyset\}$ is closed in κ.*

Proof. Let Ψ be an infinite subset of $\Sigma(P)$, and assume, for contradiction, that $\beta = \sup \Psi$ does not belong to $\Sigma(P)$, namely $V(P) \cap T_\beta = \emptyset$. By assumption, $T_\alpha \cap V(P) \neq \emptyset$ for some $\alpha < \beta$. Choose a vertex $x \in T_\alpha \cap V(P)$. Since $\beta \notin \Sigma(P)$, we have $x \notin T_\beta$, and thus $x \in RF^\circ(T_\beta)$, which together with the assumption that $\Sigma(P) \cap T_\beta = \emptyset$ implies that $V(P) \subseteq RF^\circ(T_\beta)$, meaning that $P \in \mathcal{IE}(\mathcal{Y}_\beta)$. Since $V(P) \cap T_\psi \neq \emptyset$ for every $\psi \in \Psi$, for each such ψ there exists an initial segments of P belonging to $\mathcal{E}(\mathcal{Y}_\psi)$. But this clearly implies that $P \notin \bigcup_{\psi \in \Psi} \mathcal{IE}(\mathcal{Y}_\psi)$, and thus $\beta \in \Phi_h(\mathcal{L})$, contradicting the fact that $\Phi(\mathcal{L}) \cap \Sigma = \emptyset$. □

Theorem 5.1 will follow from the combination of two theorems:

Theorem 7.29. *If Γ does not possess a hindrance or a κ-hindrance for any uncountable regular cardinal κ, then it is linkable.*

Theorem 7.30. *If Γ contains a κ-hindrance for some uncountable regular cardinal κ, then it contains a hindrance.*

Theorem 7.29 is akin to a version of the infinite "marriage theorem", proved in [11], hence an appropriate name for it is "the linkability theorem". We shall prove Theorem 7.30 in the next section, and Theorem 7.29 in the last section of the paper.

8. From κ-hindrances to hindrances

In this section we prove Theorem 7.30. Namely, that if Γ contains a κ-hindrance for some uncountable regular cardinal κ, then it is hindered. This was, in fact, proved in [8]. The proof there is only for $\kappa = \aleph_1$, but it goes verbatim to all uncountable regular cardinals κ. That proof is shorter than the one given below, since it relies on previous results. It uses the bipartite conversion, applies the bipartite version of Theorem 7.30 proved in [2], and shows how to take care of the one problem that may arise along this route, namely that the paths in the resulting hindrance are non-starting.

Our proof here does not use the main result of [2], but rather re-proves it, borrowing as "black boxes" only two lemmas. We use this as an opportunity to give the main theorem of [2] a more transparent proof, in that its main idea is summarized in a separate theorem (Theorem 8.4 below). Another advantage of the present proof is that one can see what is happening in the graph itself, rather than in the bipartite conversion.

The basic notion in the proof of the theorem is that of popularity of vertices in a hindrance. A vertex is "popular" if it has a large in-fan of $\mathcal{Y}$-alternating paths, where $\mathcal{Y}$ is the warp appearing in the hindrance, and "large" means reaching "stationarily many" points x_α. Let us first illustrate this idea in a very simple case - the simplest type of unlinkable webs:

Theorem 8.1. *A bipartite web (D, A, B) in which $|A| > |B|$ contains a hindrance.*

Proof. The argument is easy when B is finite, so assume that B is infinite, and write $|B| = \kappa$. Call a vertex $b \in B$ *popular* if $|N(b)| > \kappa$. Let U be the set of unpopular elements of B. Then $|N(U)| \leq \kappa$, and hence in the web $(D - U - N(U), A \setminus N(U), B \setminus U)$ every vertex in $B \setminus U$ is of degree larger than κ, while of course $|B \setminus U| \leq \kappa$. Hence there exists a matching F of $B \setminus U$ properly into $A \setminus N(U)$. The warp $F \cup \{(a) \mid \ a \in N(U)\}$ is then a hindrance in Δ. $\square$

Next we introduce a more general type of unlinkable webs:

Definition 8.2. A web (G, X, Y) is called κ-*unbalanced* if there exist a function $f: \ X \to \kappa$ and an injection $g: \ Y \to \kappa$, such that:

(1) $f[X]$ is κ-stationary.
(2) $f(in(P)) > g(ter(P))$ for every X–Y-path P.

This is an ordinal version of the notion of a web in which the source side has larger cardinality than the destination side. And indeed, from Fodor's lemma there follows:

Lemma 8.3. *A κ-unbalanced web is unlinkable. In fact, for every X–Y-warp $\mathcal{W}$, $f[in[\mathcal{W}]]$ is non-stationary.*

In particular, $f[X \cap Y]$ is non-stationary.

The core of the proof of Theorem 7.30 is in showing that κ-unbalanced webs are hindered, which is of course a special case of our main theorem, Theorem 5.4. But we shall need a bit more.

Given such a web, a set S of vertices is called *popular* if either $S \cap X \neq \emptyset$, or there exists an S-joined family of X-S-paths $\mathcal{P}$, such that $f[in[\mathcal{P}]]$ is κ-stationary. It is called *strongly popular* if there exists an X-S-warp $\mathcal{P}$, such that $f[in[\mathcal{P}]]$ is κ-stationary (in particular, if $f[X \cap S]$ is stationary). A vertex v is called "popular" if $\{v\}$ is popular.

Theorem 8.4. *Let $\Lambda = (G, X, Y)$ be a κ-unbalanced web, with f and g as above. Then there exists an X–Y-separating set S such that:*

(1) *Every vertex s of S is popular in $\Lambda[RF^\circ(S) \cup \{s\}]$, i.e., either $s \in X$ or there exits an X-starting s-in-fan $\mathcal{P}$ in $G[RF^\circ(S) \cup \{s\}]$, where $f[in[\mathcal{P}]]$ is stationary.*

(2) *S is not strongly popular.*
(3) $|S \setminus X| \leq \kappa$.

For the proof we shall need two results from [2]:

Lemma 8.5. *If $\Xi_u,\ u \in U$ are non-stationary subsets of κ whose union is stationary, then there exists a choice $g(u)$ of one ordinal from each Ξ_u such that $g[U]$ is stationary.*

Lemma 8.6. *With the notation above, let C be a set of vertices satisfying $|C| > \kappa$ and let $\mathcal{F}_v$ be an X-v fan for every $v \in C$. Then there exists an X–C-warp $\mathcal{F}$ such that $in[\mathcal{F}] \supseteq in[\mathcal{F}_v]$ for some $v \in C$.*

Remark: As noted in [2], Lemma 8.6 follows easily from Theorem 1.6 (assuming it is proved). In fact, Theorem 1.6 has the following stronger corollary (written below in terms of the reverse web):

Corollary 8.7 (of Theorem 1.6)**.** *Assume that the web $\Gamma = (G, A, B)$ is unlinkable, and let $\mathcal{F}_a$ be an a-B-fan for every $a \in A$. Then there exists an A–B-warp $\mathcal{F}$ such that $ter[\mathcal{F}] \supseteq ter[\mathcal{F}_a]$ for some $a \in A$.*

Proof of Corollary 8.7 Assuming the validity of Theorem 1.6, there exist a family $\mathcal{P}$ of disjoint paths and an A–B-separating set S such that S consists of a choice of one vertex from each $P \in \mathcal{P}$. Since, by assumption, Γ is unlinkable, there exists $a \in A \setminus in[\mathcal{P}]$. Then $\mathcal{P}[RF(S)]^\frown \mathcal{F}_a$ is the desired warp $\mathcal{F}$. ☐

Proof of Theorem 8.4 Let POP be the set of popular vertices of Λ, and let $UNP = V \setminus POP$. Let $U_0 = Y \cap UNP,\ P_0 = Y \cap POP$. Define inductively sets $U_i, P_i\ (i < \omega)$ as follows: $U_{i+1} = N^-(U_i) \cap UNP,\ P_{i+1} = N^-(U_i) \cap POP$. Finally, let $S = \bigcup_{i<\omega} P_i$.

Since $X \subseteq POP$, we have $U_i \cap X = \emptyset$. Let P be an X–Y-path having k vertices. By the definition of the sets U_i, if P avoids S, then $V(P) \subseteq \bigcup_{i<k} U_i$, thus $in(P) \notin X$, a contradiction. This shows that S is separating.

Assertion 8.8. *U_i is unpopular.*

Proof. By induction on i. Suppose, first, that U_0 is popular. Let $\mathcal{F}$ be a U_0-joined family of X-U_0-paths, such that $f[in[\mathcal{F}]]$ is stationary. For every $u \in U_0$ write $\mathcal{F}_u = \{P \in \mathcal{F},\ ter(P) = u\}$. For every $\alpha \in f[in[\mathcal{F}]]$ choose a path $P \in \mathcal{F}$ such that $f(in(P)) = \alpha$, and define $h(\alpha) = g(ter(P))$ (since $ter(P) \in U_0 \subseteq Y$, the value $g(ter(P))$ is defined). By Definition 8.2(2), h is regressive. Hence, by Fodor's lemma (Theorem 7.1) there exist a stationary subset Ψ of $f[in[\mathcal{F}]]$ and an ordinal β such that $h(\alpha) = \beta$ for every $\alpha \in \Psi$. This means that there exists a vertex $u \in U_0$ such that $f[in[\mathcal{F}_u]]$ is stationary, contradicting the fact that $U_0 \subseteq UNP$.

Let now $k > 0$, assume that the assertion is true for $i = k-1$, and assume, for contradiction, that U_k is popular. Let $\mathcal{F}$ be a U_k-joined family of X-U_k-paths, such that $f[in[\mathcal{F}]]$ is stationary. Again, for every $u \in U_k$ write $\mathcal{F}_u = \{P \in \mathcal{F},\ ter(P) = u\}$, and $\Xi_u = f[in[\mathcal{F}_u]]$. Since $U_k \subseteq UNP$, each set Ξ_u is non-stationary. By Lemma 8.5, there exists a choice of a path $P(u) \in \mathcal{F}_u$ for every $u \in U_k$, such that $f[in\{P(u) \mid\ u \in U_k\}]$ is stationary. Since $U_k \subseteq N^-(U_{k-1})$, by adding edges joining U_k to U_{k-1}, the family $\{P(u) :\ u \in U_k\}$ can be extended to a U_{k-1}-joined family of paths. But this contradicts the fact that U_{k-1} is unpopular. ☐

Assertion 8.9. *P_i is not strongly popular, for any $i < \omega$.*

Proof. Assume that there exists an X-P_i-warp $\mathcal{P}$ with $f[in[\mathcal{P}]]$ stationary (this happens, in particular, if $f[P_i \cap X]$ is stationary). The case $i = 0$ follows from Lemma 8.3, since $P_0 \subseteq Y$. For $i > 0$, since $P_i \subseteq N^-(U_{i-1})$, the warp $\mathcal{P}$ can be extended to a U_{i-1}-joined family of paths $\mathcal{F}$, with $in[\mathcal{F}] = in[\mathcal{P}]$. This contradicts Assertion 8.8. □

Assertion 8.10. *$|P_i \setminus X| \leq \kappa$ for every $i < \omega$.*

Proof. Every point $p \in P_i \setminus X$ has a p-joined X-p warp $\mathcal{W}_p$ such that $f(in[\mathcal{W}_p])$ is stationary. If $|P_i \setminus X| > \kappa$ then by Assertion 8.6 there exists an X-P_i-warp $\mathcal{W}$ such that $in[\mathcal{W}] \supseteq in[\mathcal{W}_p]$ for some $p \in P_i$, implying that $in[\mathcal{W}]$ is stationary, and hence that P_i is strongly popular. This contradicts Assertion 8.9. □

We are now ready to conclude the proof of Theorem 8.4. Assertion 8.10 yields condition (3) of the theorem, and Assertion 8.9 implies condition (2). It remains to show condition (1), namely that a point $s \in S$ is not only popular in Λ, but also in $\Lambda[RF^\circ(S) \cup \{s\}]$. If $s \in X$ then there is nothing to prove. Otherwise, there exists an s-joined family $\mathcal{F}$ of X-s-paths such that $f[in[\mathcal{F}]]$ is stationary. For each i let $\mathcal{F}_i$ be the set of those paths $P \in \mathcal{F}$ on which there exists a vertex $x \neq s$ in P_i such that xP meets S only at x. Since no P_i is strongly popular, $f[in[\mathcal{F}_i]]$ is non-stationary for every $i < \omega$. Hence, by Lemma 7.2, $f[in[\bigcup_{i<\omega} \mathcal{F}_i]]$ is non-stationary. Thus the set $\mathcal{F}'$ of paths from $\mathcal{F}$ meeting S only at s satisfies the property that $f[in[\mathcal{F}']]$ is stationary. □

Clearly, the properties of the set S in Theorem 8.4 imply that S is linkable in $\overleftarrow{G}$ properly into X, which yields Theorem 5.4 for κ-unbalanced webs.

Proof of Theorem 7.30.

By assumption, there exists in Γ a κ-hindrance $\mathcal{L}$ for some regular cardinal κ. We shall use for $\mathcal{L}$ the notation of Section 7. By Lemma 7.22, we may assume that $\Phi_G = \Phi_G(\mathcal{L})$ is stationary.

Let $\mathcal{Y} = \mathcal{Y}(\mathcal{L})$. We wish to turn $\mathcal{Y}$ into a hindrance. In fact, it almost *is* a hindrance: $ter[\mathcal{Y}]$ is A–B-separating, and any $\alpha \in \Phi = \Phi(\mathcal{L})$ gives rise to a path in $\mathcal{IE}(\mathcal{Y})$. The problem is that there are paths in $\mathcal{Y}$ that "hang in air", namely they start at vertices y_β. We wish to "ground" such paths, using reverse $\mathcal{Y}_G$-alternating paths from such vertices y_β to some x_α, $\alpha \in \Phi_G \setminus \Phi^\infty$ or to some infinite path H_α, $\alpha \in \Phi_G \cap \Phi^\infty$. Applying such a path to $\mathcal{Y}$ "connects y_β to the ground". We shall be able to do this only for "popular" vertices y_β, in a sense to be defined below. But using Theorem 8.4, we shall find that this suffices.

For every $\alpha \in \Phi_G \cap \Phi^\infty(\mathcal{L})$ let x_α be a new vertex added, which represents the infinite path H_α. Let X^∞ be the set of vertices thus added. Let $X = X^{fin}(\mathcal{L}) \cup X^\infty$ and $Y = Y(\mathcal{L}) \cap V[\mathcal{E}(\mathcal{Y})]$ (see Notation 7.9 for the definitions of $X^{fin}(\mathcal{L})$ and of $Y = Y(\mathcal{L})$.) To understand the choice of the definition of Y, note that only paths in $\mathcal{E}(\mathcal{Y})$ need to be "connected to the ground", to obtain a wave. For each $\alpha \leq \kappa$ write $T_\alpha = T_\alpha(\mathcal{L})$. Write $T = T_\kappa$, namely $T = ter[\mathcal{E}(\mathcal{Y})]$.

Let $\tilde{D} = D[RF(T)]$. Let F be the graph whose vertex set is $RF(T) \cup X^\infty$, and whose edge set is $E(\tilde{D}) \cup \{(x_\alpha, v) \mid u \in RF(T),\ x_\alpha \in X^\infty,\ (u, v) \in E(D) \text{ for some } u \in V(H_\alpha)\}$. Let Θ be the web (F, X, Y), and let $\Lambda = \Lambda_\Theta(\mathcal{Y})$, as defined in Section 4.2. As recalled, Λ is the web of $\mathcal{Y}$-alternating paths in Θ.

Remark 8.11. For the sake of clarity, we shall redefine the web Λ explicitly. The definition of Λ below is quite complex. However, it is quite natural when viewed in the bipartite conversion of Θ, and it is advisable to keep in mind this conversion. For example, it is helpful to remember that X consists in the bipartite conversion of "men", and hence can be connected only to "women". Since every edge $(u, v) \in E[\mathcal{Y}]$ corresponds to the edge $(m(u), w(v))$ in the bipartite conversion, this means that $x \in X$ can be connected in Λ only to v.

The vertex set of Λ is $V_\Lambda = X \cup Y \cup (RF(T) \setminus V[\mathcal{Y}]) \cup E[\mathcal{Y}]$.

The edge set of Λ is constructed by the rule that an edge $(u, v) \in E[\mathcal{Y}]$ sends an edge somewhere (namely, a vertex or an edge) if u sends there an edge in D, and it receives an edge from somewhere if v receives an edge from there (corresponding to an edge ending at $w(v)$). We shall also have edges between two consecutive edges (u, v) and (v, w) of $\mathcal{Y}$, the edge being directed from the latter to the former (in the bipartite conversion this means "directed from the man to the woman". In alternating paths terminology, this corresponds to the fact that alternating paths go backwards on paths from $\mathcal{Y}$). Another rule is that X-vertices only send edges, and Y vertices only receive edges. Finally, a vertex $x_\alpha \in X^\infty$ sends edges in Λ to all vertices (and, consequently, to edges) to which some vertex on H_α sent an edge in D.

Formally, write:

$$E_{VV} = \{(u, v) \mid \ u \in (RF(T) \setminus V[\mathcal{Y}]) \cup X,\ v \in (RF(T) \setminus V[\mathcal{Y}]) \cup Y,\ (u, v) \in E(D)\}$$
$$E_{EV} = \{(e, w) \mid e = (u, v) \in E[\mathcal{Y}],\ w \in (RF(T) \setminus V[\mathcal{Y}]) \cup Y,\ (u, w) \in E(D)\}$$
$$E_{VE} = \{(w, e) \mid e = (u, v) \in E[\mathcal{Y}],\ w \in (RF(T) \setminus V[\mathcal{Y}]) \cup X,\ (w, v) \in E(D)\}$$
$$E_{EE} = \{(e, f) \mid e = (u, v),\ f = (w, z) \in E[\mathcal{Y}], u = z \text{ or } (v, w) \in E(D)\}$$
$$E_{\infty V} = \{(x_\alpha, v) \mid x_\alpha \in X^\infty,\ v \in (RF(T) \setminus V[\mathcal{Y}]) \cup Y, (u, v) \in E(D) \text{ for some } u \in H_\alpha\}$$
$$E_{\infty E} = \{(x_\alpha, e) \mid x_\alpha \in X^\infty,\ e = (w, v) \in E[\mathcal{Y}],\ (u, v) \in E(D) \text{ for some } u \in H_\alpha\}$$

Let $E_\Lambda = E_{VV} \cup E_{EV} \cup E_{VE} \cup E_{EE} \cup E_{\infty V} \cup E_{\infty E}$. Let D_Λ be the digraph (V_Λ, E_Λ), and define the web Λ as (D_Λ, X, Y). For each $x = x_\alpha \in X$ define $f(x) = \alpha$, and for each $y = y_\beta \in Y$ let $g(y) = \beta$.

Assertion 8.12. *Λ is κ-unbalanced, as is witnessed by f and g.*

Proof. Condition (1) of Definition 8.2 is true since $f[X] = \Phi(\mathcal{L})$. Condition (2) is tantamount to the fact that $g(ter(Q)) < f(in(Q))$ for every X–Y-alternating path Q in Θ. If $in(Q) \in X^{fin}$ then this follows from Lemmas 7.17 and 7.18. If $in(Q) = x_\alpha \in X^\infty$, and the first edge in Q is (x_α, u), then in D there exists an edge (v, u) for some $v \in H_\alpha$. Then $v \in RF(T_\gamma)$ for some $\gamma \leq \alpha$, and thus, again by Lemma 7.18, $g(ter(Q)) < \gamma$, yielding $g(ter(Q)) < \alpha$. □

Let S be an X–Y-separating set as in Theorem 8.4. Write $S_V = S \cap V(D)$, $S_E = S \cap E[\mathcal{Y}]$. Also write $\Theta - S$ for the web obtained from Θ by deleting S_V from its vertex set, and deleting S_E from its edge set.

The fact that S is X–Y-separating in Λ implies that there are no augmenting $\mathcal{Y}$-alternating paths in $\Theta - S$. Namely:

Assertion 8.13. *There are no S-avoiding $\mathcal{Y}$-alternating paths in D from X to Y.*

Let $\mathcal{G} = \mathcal{Y} - S_E$, namely the set of fragments of $\mathcal{Y}$ resulting from the deletion of edges in S_E.

Remark 8.14. To understand the next assertion, it should be kept in mind that there are $\mathcal{Y}$-alternating paths that start at some x_α, and as their first step go backwards on an edge belonging to $E[\mathcal{Y}]$. This type of alternating paths is again best understood in terms of the bipartite conversion. In the bipartite conversion, the first edge of the corresponding alternating path starts with the edge $(m(x_\alpha), w(x_\alpha))$, which does not belong to $E[\mathcal{Y}]$, as is the customary definition of alternating paths.

Assertion 8.15. *Let $H = H_\alpha$ be a path belonging to $\mathcal{G}^f_G$ (H is then a finite path in $\mathcal{IE}(\mathcal{Y})$ not containing an edge from S_E), such that $x = ter(H) \in X \setminus S$. Then there is no $\mathcal{Y}$-alternating path avoiding S from a vertex of H to $Y \setminus S$.*

Proof. Suppose that there exists such a path Q. Let u be the last vertex on Q lying on H. Then the path $\overleftarrow{H}uQ$ is a $\mathcal{Y}$-alternating X–Y-path avoiding S (see the remark above), contradicting the fact that S is separating in Λ. □

Notation 8.16. Denote by $\mathcal{H}_\emptyset$ the set of paths $H = H_\alpha \in \mathcal{G}_G$ such that either:

(i) H is finite and $ter(H) \notin S$, or:

(ii) H is infinite and no $\mathcal{Y}$-alternating, S-avoiding path starts at a vertex of H and ends at $Y \setminus S$.

Let $\mathcal{G}' = \mathcal{G} \setminus \mathcal{H}_\emptyset$.

Let RR be the set of vertices v such that there exists an S-avoiding $\mathcal{G}$-alternating path starting at v and terminating at $Y \setminus S$. Assertion 8.15 implies:

Assertion 8.17. *If $P \in \mathcal{G}$ and $V(P) \cap RR \neq \emptyset$ then $P \in \mathcal{G}'$.*

For each $P \in \mathcal{G}'$ define $bl(P)$ to be:

- the first vertex on P belonging to RR if $V(P) \cap RR \neq \emptyset$, and:
- $ter(P)$, if $V(P) \cap RR \neq \emptyset$.

Let $BL = \{bl(P) \mid P \in \mathcal{G}'\}$ and $BB = S_V \cup BL$.

Assertion 8.18. *BB is A–B-separating.*

(Remark: The idea of the proof is borrowed from the proof of Theorem 4.7.)

Proof. Since T is A–B-separating, it suffices to show that BB is A-T-separating. Let R be an A-T-path in D, and assume, for contradiction, that $V(R) \cap BB = \emptyset$. Write $t = ter(R)$. Since $t \in T = \mathcal{E}(ter[\mathcal{Y}])$, and since by assumption $t \notin S_V$, it follows that $t = ter(P)$ for some path $P \in \mathcal{G}$. Since P is finite, and since $ter(P) \in \mathcal{E}(ter[\mathcal{Y}])$ (namely, P cannot be some H_α), $P \in \mathcal{G}'$. Let $q = bl(P)$. Since $t \notin BB$, it follows that $t >_P q$. Let Q be a $\mathcal{G}$-alternating path from q to $Y \setminus S$.

Assume, first, that R does not meet any path of $\mathcal{G}$ apart from P. Then, in particular, $in(R) \notin V[\mathcal{Y}]$, and hence $in(R) \in X$. If R does not meet Q, then the path $Rt\overleftarrow{P}qQ$ is an S-avoiding $\mathcal{Y}$-alternating path from A to Y, contradicting Assertion 8.13. If R meets Q, and the last vertex on R belonging to Q is, say, v then RvQ is an S-avoiding $\mathcal{Y}$-alternating path from A to Y, again providing a contradiction.

Thus we may assume that R meets another path from $\mathcal{G}$, besides P. Let P_1 be the last path different from P met by R, and let t_1 be the last vertex on R lying on P_1. The path $t_1Rt\overleftarrow{P}Z$ (or a "shortcut" of it, as in the previous paragraph) witnesses the fact that $t_1 \in RR$, and hence by Assertion 8.17 $P_1 \in \mathcal{G}'$. Let $q_1 = bl(P_1)$. Since by assumption $v_1 \notin BB$, it follows that $t_1 >_{P_1} q_1$. Let Q_1 be an S-avoiding $\mathcal{G}$-alternating path from q_1 to $Y \setminus S$. If R does not meet any other path, besides P and P_1, belonging to $\mathcal{G}$ then the path $Rt_1\overleftarrow{P_1}q_1Q_1$ (or a shortcut of it) is an S-avoiding X–Y $\mathcal{G}$-alternating path, contradicting Assertion 8.13. Thus we may assume that R meets still another path from $\mathcal{G}$. Continuing this argument, we eventually must reach a contradiction, since R is finite. □

Assertion 8.19. *Let $p \in RF(T)$, and let $\mathcal{J}$ be an X-p-in-fan of $\mathcal{Y}$-alternating paths in Θ, such that each path in $\mathcal{J}$ meets some path in $\mathcal{Y}_H$ not containing p. Then $f[in[\mathcal{J}]]$ is non-stationary.*

Proof. Assume for contradiction that $f[in[\mathcal{J}]]$ is stationary. For each $P \in \mathcal{J}$ choose $\beta = \beta(P)$ such that P meets the path $\mathcal{Y}(y_\beta)$. As before, by choosing a subfamily of $\mathcal{J}$ if necessary, we may assume that f is injective on $in[\mathcal{J}]$. Hence the function h on $f[in[\mathcal{J}]]$ defined by $h(\alpha) = \beta(P)$ for that $P \in \mathcal{J}$ for which $f(in(P)) = \alpha$, is well defined. By an argument as in the proof of Assertion 8.12, $h(\alpha) < \alpha$, namely h is regressive. By Fodor's Lemma, this implies that $f^{-1}(\beta)$ is of size κ for some β. But this is clearly impossible, since only finitely many paths from $\mathcal{J}$ can meet $\mathcal{Y}(y_\beta)$. □

Assertion 8.20. *Let $p \in RF(T)$, and let $\mathcal{J}$ be an X-p-fan of $\mathcal{Y}$-alternating paths in Θ, such that each path in $\mathcal{J}$ meets a path in $\mathcal{G}_H$ (namely, a fragment of $\mathcal{Y} - S_E$ hanging in air) not containing p. Then $f[in[\mathcal{J}]]$ is non-stationary.*

Proof. Suppose that $f[in[\mathcal{J}]]$ is stationary. Let $P \in \mathcal{J}$. Choose a path $W \in \mathcal{G}_H$ that P meets, and let e be the last edge of P lying on W. Denote by s the edge in S_E such that $head(s) = in(W)$. Going from s along W to e and then continuing along P yields then a $\mathcal{Y}$ alternating path $Q(P)$ starting at s and ending at $ter(P)$. Since the paths $Q(P)$ are all disjoint, it follows that S_E is strongly popular. But this contradicts property (3) of S_E, as guaranteed by Theorem 8.4. □

Assertion 8.21. *Let Q be an X-starting $\mathcal{Y}$-alternating path avoiding S. Suppose that Q meets a path P from $\mathcal{G}$, and let p be the last point on P belonging to Q (thus $p = tail(e)$ for some edge $e \in E(P) \cap E(\overleftarrow{Q})$). Then $p \leq_P bl(P)$.*

Proof. Assume that $bl(P) <_P p$. By the definition of $bl(P)$, there exists a $\mathcal{Y}$-alternating path R, starting at $bl(P)$, ending in Y and avoiding S. Then the $\mathcal{Y}$-alternating path $Qp\overleftarrow{P}bl(P)R$ (or part of it, if R meets Q,) is an S-avoiding X–Y $\mathcal{Y}$-alternating path, contradicting the fact that S is X–Y-separating in Λ. □

Assertion 8.22. *There exists in Γ a warp $\mathcal{V}$ such that $in[\mathcal{V}] \subseteq A$ and $ter[\mathcal{V}] = BB$.*

Proof. Let $\tilde{S} = S_V \setminus X \cup \{head(e) \mid e \in S_E\}$. Order the points of $\tilde{S}$ as $(s_\theta : \theta < \lambda)$, where $\lambda \leq \kappa$. By the properties of S, each s_θ has an X-s_θ-fan $\mathcal{F}_\theta$ in $\Theta - S$ of size κ of $\mathcal{Y}$-alternating paths, such that $f[in[\mathcal{F}_\theta]]$ is stationary. By Assertion 8.19 we may also assume that no path in $\mathcal{F}_\theta$ meets a path from $\mathcal{Y}_H$, namely:

(i) All paths in $\mathcal{F}_\theta$ meet (apart from possibly at s_θ) only paths from $\mathcal{Y}_G$.

By Assertion 8.20 we may further assume that no path in $\mathcal{F}_\theta$ meets a path in $\mathcal{G}_H$, namely:

(ii) All paths in $\mathcal{F}_\theta$ meet (apart from possibly at s_θ) only paths from $\mathcal{G}_G$.

By induction on θ, choose for each s_θ a $\mathcal{Y}$-alternating path $Q_\theta \in \mathcal{F}_\theta$, ending at s_θ and satisfying:

(a) Q_θ does not meet any path from $\mathcal{Y}_G$ met by any $Q_\delta,\ \delta < \theta$.

(b) Q_θ does not meet (apart from possibly at s_θ) any path from $\mathcal{Y}_H$.

(c) Q_θ does not meet (apart from possibly at s_θ) any path from $\mathcal{G}_H$.

Since the paths Q_θ avoid S, they are not only $\mathcal{Y}$-alternating, but also $\mathcal{G}$-alternating. We now apply all Q_θ's to $\mathcal{G}$. Let $\mathcal{Z}$ be the resulting warp. We wish to form a corresponding warp in D. The paths in $\mathcal{Z}$ which are not contained in D are paths Z such that $in(Z) = x_\alpha \in X^\infty$. Such a path was obtained by the application of an alternating path Q_θ such that $in(Q_\theta) = x_\alpha$. Let (x, v) be the first edge of Q_θ. By the definition of $E(\Lambda)$, this means that $(p, v) \in E(D)$ for some $p \in V(H_\alpha)$. Replace then Z by $H_\alpha p Z$.

Denote by $\mathcal{U}$ the resulting warp in D. Conditions (a), (b) and (c) imply that there are no non-starting paths in $\mathcal{U}$ and $in[\mathcal{U}] \subseteq A$. Assertion 8.21 together with condition (a) imply that each path from $\mathcal{U}$ intersects BB at most once. Assertion 8.21 also implies $BB \subseteq V[\mathcal{U}]$. Therefore, by pruning the warp $\mathcal{U}$ we can obtain a warp $\mathcal{V}$ with $in[\mathcal{V}] \subseteq A$ and $ter[\mathcal{V}] = BB$ as required.

$\square$

Since BB is separating, $\mathcal{V}$ is a wave. By the equivalent formulation of the main theorem, given in Conjecture 5.2, to complete the proof of the theorem it is enough to show that $\mathcal{V}$ is non-trivial, which is clear. In fact, more than that is true: $\mathcal{E}(\mathcal{V})$ is a hindrance, in a strong sense. Since S is not strongly popular in Λ, the set $\{f(ter(Q_\theta) \mid \theta < \lambda\}$ is non-stationary. Thus, the set $\Xi = \{\alpha \mid x_\alpha \notin ter[\mathcal{V}]\}$ is stationary. Each $\alpha \in \Xi$ corresponds to some (finite or infinite) path H_α, unreached by any Q_θ, and thus belonging to $\mathcal{IE}(\mathcal{V})$.

This completes the proof of Theorem 7.30. To prove Theorem 5.4, and thereby Theorem 1.6, it remains to prove the "linkability theorem", Theorem 7.29.

9. Proof of the Linkability Theorem

Define the *height* of a set Y of vertices to be the minimal cardinality of a subset X of $V \setminus A$ for which there exists a wave $\mathcal{W}$ in Γ/X, such that $Y \subseteq RF_\Gamma(ter[\mathcal{W}])$. The height of Γ is defined as the height of V.

Definition 9.1. A warp $\mathcal{W}$ is a *half-way linkage* if it is an A–C-linkage, with $ter[\mathcal{U}] \subseteq C$, for some minimal separating set C for which Γ/C is unhindered. Such a set C is called a *stop-over set* of $\mathcal{W}$. Note that in this definition C is not uniquely determined by $\mathcal{W}$. The *altitude* of $\mathcal{W}$ is the minimal height of such a set C.

We shall prove:

Theorem 9.2. *Suppose that Γ is unhindered. Let $A' \subseteq A$ be a set of cardinality λ. Then*

- *(♣) If $(D, A \setminus A', B)$ is linkable then so is the web (D, A, B).*

- (♣♣) *There exists a half-way linkage of altitude at most λ, linking A' to B.*

Theorem 7.29 follows from (♣) upon taking $A' = A$.

To gradually impart the ideas of the proof of Theorem 9.2, let us first prove a few low cardinality cases.

Proof of (♣) for $\lambda = \aleph_0$. This is the main result of [6]. The proof there is very laborious, circumventing as it does Theorem 6.1. With the aid of the latter, (♣) follows in the countable case by a classic "Hilbert hotel" argument. Let $\mathcal{F}$ be a linkage in the web $(D, A \setminus A', B)$. Let $A_0 = A'$. Choose a vertex $a \in A_0$, and using Theorem 6.1 link it to B by a path P_1, such that $\Gamma - P_1$ is unhindered. Let $A_1 = A_0 \cup in[\mathcal{F}\langle V(P_1)\rangle]$ (namely, A_1 is obtained by adding to A_0 all initial points of paths from $\mathcal{F}$ met by P_1). Choose a vertex from A_1, different from a, and link it to B by a path P_2 in $\Gamma - P_1$, such that $\Gamma - P_1 - P_2$ is unhindered. Let $A_2 = A_1 \cup in[\mathcal{F}\langle V(P_2)\rangle]$. Continuing this way, and choosing wisely the order of the elements to be linked by P_i, all elements of all A_i's serve as $in(P_j)$ for some j, and thus the set $A'' = \bigcup A_i$ is linked to B by the warp $\mathcal{P} = \{P_0, P_1, \ldots\}$, and all paths in $\mathcal{F}\langle A \setminus A''\rangle$ are disjoint from all paths in $\mathcal{P}$. Thus $\mathcal{F}\langle A \setminus A''\rangle \cup \mathcal{P}$ is a linkage of A.

Proof of (♣♣) for $\lambda = \aleph_0$ and $|V| = \aleph_1$. Order the elements of V as $(v_\theta : \theta < \aleph_1)$. Construct an $\aleph_1$-ladder $\mathcal{L}$, at each stage α choosing y_α to be the first v_θ not belonging to $RF(T_\alpha)$ and choosing $\mathcal{W}_\alpha$ to be a hindrance in Γ_α if such exists. The construction of $\mathcal{L}$ terminates after $\zeta \leq \aleph_1$ steps.

By the choice of the vertices y_α, we have:

Assertion 9.3. $V = \bigcup_{\alpha \in \Sigma} RF(T_\alpha) = RF(\mathcal{L})$.

Write $\mathcal{Y} = \mathcal{Y}(\mathcal{L})$ and for $\alpha \leq \zeta$ write $\mathcal{Y}_\alpha = \mathcal{Y}_\alpha(\mathcal{L})$ (thus $\mathcal{Y} = \mathcal{Y}_\zeta$) and $T_\alpha = T_\alpha(\mathcal{L})$.

Assume, first, that ζ is countable. By Assertion 9.3 $RF[T_\zeta] = V$ and hence $T_\zeta = \mathcal{E}(V) = B$. Together with Lemma 7.26 (applied with $\alpha = \zeta$) this implies that $\mathcal{Y}\langle\sim B\rangle$ is countable. Thus, $A \setminus in[\mathcal{Y}\langle B\rangle]$ is countable. Hence, by the case of (♣) proved above, Γ is linkable, which clearly implies (♣♣).

Thus we may assume that $\zeta = \aleph_1$. Since Γ is unhindered, by Theorem 7.30 $\mathcal{L}$ is not an $\aleph_1$-hindrance, and hence there exists a closed unbounded set Σ not intersecting $\Phi(\mathcal{L})$. By Lemma 7.6, $\Sigma \cap \Phi_h(\mathcal{L}) = \emptyset$, namely:

Assertion 9.4. *Γ_α is unhindered for every $\alpha \in \Sigma$.*

Assertion 9.3 implies:

Assertion 9.5. *For every countable set of vertices X there exists $\gamma(X) \in \Sigma$ such that $X \subseteq RF(T_{\gamma(X)})$.*

Assertion 9.6. *$\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle T_\beta\rangle$ is countable for every $\alpha, \beta \in \Sigma$.*

Proof. If $\beta < \alpha$ then $\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle T_\beta\rangle$ consists of those paths in $\mathcal{Y}$ that start at some y_γ for some $\beta \leq \gamma < \alpha$, and thus it is countable. For $\alpha < \beta$, we have $\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle T_\beta\rangle \subseteq \mathcal{IE}(\mathcal{Y}_\beta)$, and hence the assertion follows from Lemma 7.8. □

In particular, $\mathcal{Y}_G \setminus \mathcal{Y}\langle T_\alpha\rangle = \mathcal{Y}\langle T_0\rangle \setminus \mathcal{Y}\langle T_\alpha\rangle$ is countable for every $\alpha \in \Sigma$ (remember that "$\mathcal{Y}_G$" stands for "$\mathcal{Y}\langle A\rangle$").

Write $A_0 = A'$. Choose $a_0 \in A_0$, and using Theorem 6.1 link it to B by a path P_0, such that $\Gamma - P_0$ is unhindered. Let $\gamma_0 = \gamma(V(P_0))$. (See Assertion 9.5 for the definition of γ.) Let $A_1 = A_0 \cup in[\mathcal{Y}_G\langle V(P_0)\rangle] \cup in[\mathcal{Y}_G \setminus \mathcal{Y}\langle T_{\gamma_0}\rangle]$. By Assertion 9.6 A_1 is countable.

Choose $a_1 \in A_1 \setminus \{a_0\}$, and find an a_1-B path P_1 such that $\Gamma - P_0 - P_1$ is unhindered. Let $\gamma_1 = \max(\gamma(V(P_0)), \gamma(V(P_1)))$, and $A_2 = A_1 \cup in[\mathcal{Y}_G\langle V(P_1)\rangle] \cup in[\mathcal{Y}_G \setminus \mathcal{Y}\langle T_{\gamma_1}\rangle]$.

Continue this way ω steps. Let $X = \cup_{i<\omega} V(P_i)$, and $\gamma = \sup_{i<\omega} \gamma_i$. Since Σ is closed, $\gamma \in \Sigma$. By Lemma 7.28 every path $P \in \mathcal{Y}_G \setminus \mathcal{Y}\langle T_\gamma\rangle$ must belong to $\mathcal{Y}_G \setminus \mathcal{Y}\langle T_{\gamma_i}\rangle$ for some $i < \omega$ and then, by the definition of the sets A_i, we have $in(P) \in A_{i+1}$. Note that each path P_i ends at some vertex in $B \cap RF(T_\gamma)$ and since a vertex in B can only be roofed by itself, this vertex must be in T_γ.

Choosing the vertices a_i in an appropriate order, we can see to it that $\{a_i : i < \omega\} = A' \cup in[\mathcal{Y}_G \setminus \mathcal{Y}\langle T_\gamma\rangle] \cup in[\mathcal{Y}\langle X\rangle]$. Write $\mathcal{P} = \{P_i : i < \omega\}$, and let $\mathcal{V} = \mathcal{P} \cup \mathcal{Y}\langle \sim X\rangle[RF(T_\gamma)]\langle A\rangle$. Then $\mathcal{V}$ is an A-T_γ-linkage linking A' to B. By Assertion 9.4, Γ/T_γ is unhindered and therefore, taking $C = T_\gamma$ in the definition of "half-way linkage" shows, by Lemma 7.11, that $\mathcal{V}$ is a half-way linkage. The warp $\mathcal{Y}_\gamma/Y_\gamma(\mathcal{L})$ is a wave in $\Gamma/Y_\gamma(\mathcal{L})$, whose terminal points set contains T_γ, showing that $\mathcal{V}$ has countable altitude.

This concludes the proof of (♣♣) for $\lambda = \aleph_0$ and $|V| = \aleph_1$.

Proof of (♣) for $\lambda = |V| = \aleph_1$. This was proved in [8], assuming Theorem 6.1. The arguments given here are more involved, but fit better our general proof scheme.

We may clearly assume that $A' = A$. Again, construct an $\aleph_1$-ladder $\mathcal{L}$, for which Assertion 9.3 holds. Let Σ be defined as above (once again using Theorem 7.30).

In the construction of $\mathcal{L}$, we take each $\mathcal{W}_\alpha$ to be a hindrance in Γ_α, if such exists. By Corollary 3.21, we may also assume that $\mathcal{W}_\alpha$ is a maximal wave in Γ_α ($\vec{\preccurlyeq}$-maximal and thus also $\leq$-maximal). The maximality of $\mathcal{W}_\alpha$ implies:

Assertion 9.7. *For all $\alpha < \aleph_1$, every wave in Γ_α is roofed by $T_{\alpha+1}$.*

which implies:

Corollary 9.8. *Whenever $\alpha < \beta < \aleph_1$, every wave in Γ_α is roofed by T_β.*

Assertion 9.9. *If $\alpha < \zeta$ and $X \subseteq RF_{\Gamma_\alpha}(T_\zeta)$ then every wave in Γ_α/X is roofed by $T_{\zeta+1}$.*

Proof. Let $\mathcal{V}$ be a wave in Γ_α/X. Then $\mathcal{V}/T_\zeta$ is a wave in $(\Gamma_\alpha/X)/T_\zeta = \Gamma_\zeta$. By Corollary 9.8, the wave $\mathcal{V}/T_\zeta$ is roofed by $T_{\zeta+1}$, which implies that $\mathcal{V}$ is roofed by $T_{\zeta+1}$. □

The core of the proof is in the following:

Assertion 9.10. *Let α be an ordinal in Σ, and let U be a countable subset of T_α. Then there exist $\beta > \alpha$ in Σ and a T_α-T_β linkage $\mathcal{T}$ linking U to B, such that all but at most countably many paths of $\mathcal{T}$ are contained in paths of $\mathcal{Y}$.*

Proof. By the special case of (♣♣) proved above, there exists in Γ_α a half-way linkage $\mathcal{U}$ of altitude $\aleph_0$, linking U to B. Let C be a stop-over set of $\mathcal{U}$, of height $\aleph_0$. We claim that there exists $\beta > \alpha$ in Σ such that $C \subseteq RF(T_\beta)$. The fact that $\mathcal{U}$ has altitude $\aleph_0$ means that C is roofed by a wave in $(\Gamma/T_\alpha)/X$ for some countable set X. Take $\beta \in \Sigma$ such that $\beta > \max(\alpha, \gamma(X))$ (where $\gamma(X)$ is defined as in Assertion 9.5, which is valid also in the present case). By Assertion 9.9 we know that every wave in $(\Gamma/T_\alpha)/X$ is roofed by T_β and thus also C is roofed by T_β.

By Lemma 2.19, the set C is T_α–T_β-separating, and thus

$$\mathcal{Y}\langle T_\alpha\rangle\langle T_\beta\rangle \subseteq \mathcal{Y}\langle C\rangle. \tag{4}$$

Note that Assertion 9.6 holds here (with the same proof as in the previous case), and together with Equation (4), it yields:

$$|\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle C\rangle| \leq \aleph_0. \tag{5}$$

Let J be the graph on $V(D)$ whose edge set is $E[\mathcal{U}] \cup E[\mathcal{Y}]$. By (5), at most countably many connected components of J contain vertices of U or paths from $\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle C\rangle$. In all other connected component of J we can replace the paths of $\mathcal{U}$ by the segments of the paths of $\mathcal{Y}$ between T_α and C while maintaining the properties of $\mathcal{U}$ as being a T_α-C linkage linking U to B. Therefore we may assume that all but countably many paths in $\mathcal{U}$ are contained in paths of $\mathcal{Y}$.

Similarly to (5) we have:

$$|\mathcal{Y}\langle T_\beta\rangle \setminus \mathcal{Y}\langle C\rangle| \leq \aleph_0. \tag{6}$$

This implies that there exists a warp $\mathcal{F}$, whose paths are parts of paths of $\mathcal{Y}$, linking all but countably many vertices of $ter[\mathcal{U}]$ to T_β.

We may clearly assume (and hence will assume) that each path $P \in \mathcal{U}$ meets C only at $ter(P)$ and therefore $V[\mathcal{U}] \setminus ter[\mathcal{U}] \subseteq RF^\circ(C)$. However, a path $F \in \mathcal{F}$ may intersect C many times and may pass through $RF^\circ(C)$. We wish to use F in the construction of the desired linkage $\mathcal{T}$, which explains the necessity of the term $V[\mathcal{F}]$ in the following definition: define Δ as the web $(D[(RF(T_\beta) \setminus RF^\circ(C)) \cup V[\mathcal{F}]], ter[\mathcal{U}], T_\beta)$. Clearly, $\Delta/(C \setminus ter[\mathcal{U}]) = (\Gamma/C)[RF(T_\beta)]$. By Observation 3.2 since Γ/C is unhindered so is $\Delta/(C \setminus ter[\mathcal{U}])$, and hence by Corollary 3.6 Δ is unhindered.

We now apply the case $\lambda = \aleph_0$ of (♣) to Δ and $A' = ter[\mathcal{U}] \setminus in[\mathcal{F}]$. This gives a linkage $\mathcal{Q}$ of $ter[\mathcal{U}]$ to T_β. By arguments similar to those given above, we may assume that all but countably many paths of $\mathcal{Q}$ are contained in paths of $\mathcal{Y}$. The concatenation $\mathcal{U} * \mathcal{Q}$ is then the linkage $\mathcal{T}$ desired in the assertion. □

We now use Assertion 9.10 to prove (♣). The general idea of the proof is to link "slices" of the web, lying between T_α's, for ordinals $\alpha \in \Sigma$. Assertion 9.10 is used to avoid the generation of infinite paths in this process. By Lemma 7.7, paths belonging to $\mathcal{Y}$ do not become infinite along this procedure. Thus we have to be careful only about paths not contained in paths from $\mathcal{Y}$. Using the assertion, at each stage we can take care of such paths, by linking their terminal points to B.

Formally, this is done as follows. Write A as $\{a_\alpha : \alpha < \omega_1\}$, and let $U_0 = \{a_0\}$. Use the assertion to find $\sigma_1 < \omega_1$ in Σ and an A-T_{σ_1} linkage $\mathcal{T}_0$, linking a_0 to B, such that at most countably many paths of $\mathcal{T}_0$ are not contained in a path of $\mathcal{Y}$. Let U_1 be the set of end vertices of such paths, together with the end vertex of the path in $\mathcal{T}_0$ starting at a_1.

We use the assertion in this way, to define inductively ordinals $\sigma_\alpha \in \Sigma$ and T_{σ_α}-$T_{\sigma_{\alpha+1}}$ linkages $\mathcal{T}_\alpha$ linking U_α to B. Having defined these up to and including α, we write $\mathcal{T}_{\leq\alpha} = *(\mathcal{T}_\theta : \theta \leq \alpha)$ and $\mathcal{T}_{<\alpha} = *(\mathcal{T}_\theta : \theta < \alpha)$. Let $U_{\alpha+1}$ consist of the end vertices of all paths in $\mathcal{T}_{\leq\alpha}$ not contained in a path of $\mathcal{Y}$, together with the end vertex of the path in $\mathcal{T}_{\leq\alpha}$ starting at $a_{\alpha+1}$.

Assertion 9.11. *$\mathcal{T}_{<\alpha}$ is an $A - S_{\sigma_\alpha}$ linkage.*

Proof. For successor α, this follows by induction from the definitions. For limit α, this follows from Lemma 7.28, and the fact that, by our construction, all paths in $\mathcal{T}_{<\alpha}$ not contained in a path from $\mathcal{Y}$ terminate in B. □

For limit α we take $U_\alpha = ter[\mathcal{T}_{<\alpha}\langle\{a_\alpha\}\rangle]$ and $\sigma_\alpha = \sup_{\theta<\alpha}\sigma_\theta$.

Since a_α is linked to B by $\mathcal{T}_\alpha$, the concatenation $\mathcal{T}$ of $(\mathcal{T}_\alpha : \alpha < \omega_1)$ is the desired A–B linkage.

This concludes the proof of (♣) for $\lambda = |V| = \aleph_1$.

We now go on to the proof of (♣) and (♣♣) in the general case.

Proof of (♣) (assuming (♣) and (♣♣) for cardinals smaller than λ)
Case I: λ is regular.

Let $\mathcal{F}$ be a linkage in the web $(D, A \setminus A', B)$. Similarly to the $\lambda = \aleph_1$ case, we construct a λ-ladder $\mathcal{L}$ and a choose a closed unbounded set $\Sigma \subseteq \lambda$ disjoint from $\Phi(\mathcal{L})$. At each stage α we take $\mathcal{W}_\alpha$ to be a maximal hindrance in Γ_α, if Γ_α is hindered. Then Corollary 9.8 and Assertion 9.9 are valid also here.

Let $\mathcal{Y} = \mathcal{Y}(\mathcal{L})$. We then have the analogue of Assertion 9.6:

Assertion 9.12. *$|\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle T_\beta\rangle| < \lambda$ for every $\alpha, \beta \in \Sigma$.*

(For the notation used, see Convention 7.14.)

The difficulty we may face is that possibly $|V| > \lambda$. This means that Assertion 9.3 may fail, namely we cannot guarantee that every vertex is roofed by some T_α. We can only hope to achieve this for λ many vertices. Fortunately, this suffices. Along with the construction of the rungs R_α of $\mathcal{L}$, we shall define sets Z_α of cardinality at most λ, each of whose elements we shall wish to roof by T_β for some $\beta > \alpha$.

Having defined Z_θ, we enumerate its elements as $(z_\theta^\beta :\ \beta < |Z_\theta| \leq \lambda)$.

To define Z_α, we do the following. Assume that the rungs R_β of $\mathcal{L}$ as well as the sets Z_β have been defined for $\beta < \alpha$. Write $Z_{<\alpha} = \bigcup_{\theta<\alpha} Z_\theta$ and $Z_{<\alpha}^{<\alpha} = \{z_\beta^\gamma :\ \beta < \alpha,\ \gamma < \alpha\}$.

Let (γ, δ) be a pair of ordinals such that $\alpha = \max(\gamma, \delta)$. Consider two cases:

- Γ_δ is unhindered. Apply then (♣♣), which by the inductive hypothesis is true when $|A'| < \lambda$, to the web Γ_δ with $A' = T_\delta \cap Z_{<\gamma}^{<\gamma}$. This yields the existence of a half-way linkage $\mathcal{A} = \mathcal{A}_{\delta,\gamma}$ in Γ_δ, linking $T_\delta \cap Z_{<\gamma}^{<\gamma}$ to B. Furthermore, $\mathcal{A}$ is of height less than λ, namely it is roofed by some wave in $\Gamma_\delta / X_{\delta,\gamma}$ for some set $X_{\delta,\gamma}$ of cardinality less than λ.
- Γ_δ is hindered. In this case let $X_{\delta,\gamma} = \emptyset$.

Let (β, γ, δ) be a triple of ordinals such that $\delta < \beta$ and $\alpha = \max(\beta, \gamma)$. Consider the following two cases:

- There exists a T_δ-T_β-linkage linking $T_\delta \cap Z_{<\gamma}^{<\gamma}$ to B, in which all paths are contained in paths of $\mathcal{Y}_\beta$ except for a set of size smaller than λ. In such a case choose such a linkage and denote it by $\mathcal{U}_{\beta,\gamma,\delta}$. Write $\mathcal{U}_{\beta,\gamma,\delta}^m$ for the set of paths in $\mathcal{U}_{\beta,\gamma,\delta}$ not contained in a path of $\mathcal{Y}$ (the "m" standing for "maverick").

- There does not exist such a linkage. Write then $\mathcal{U}^m_{\beta,\gamma,\delta} = \emptyset$.

Let $Z_0 = A'$ and for $\alpha > 0$ let

$$Z_\alpha = Z_{<\alpha} \cup V(H_\alpha) \cup \{y_\alpha\} \cup V[\mathcal{F}\langle Z_{<\alpha}\rangle] \cup V[\mathcal{Y}\langle Z_{<\alpha}\rangle] \cup \bigcup_{\substack{\delta \leq \alpha \\ \gamma \leq \alpha}} X_{\delta,\gamma} \cup \bigcup_{\substack{\delta < \beta \leq \alpha \\ \gamma \leq \alpha}} V[\mathcal{U}^m_{\beta,\gamma,\delta}]$$

Let $Z = \bigcup_{\alpha<\lambda} Z_\alpha$. By the regularity of λ we have:

Assertion 9.13. *Every subset U of Z of cardinality less than λ is contained in $Z^{<\alpha}_{<\alpha}$ for some $\alpha < \lambda$.*

Choosing carefully the vertices y_α in the ladder $\mathcal{L}$, we can see to it that the following weaker version of Assertion 9.3 holds:

Assertion 9.14. $Z \subseteq RF(\mathcal{L})$.

We now have the analogue of Assertion 9.10, with practically the same proof:

Assertion 9.15. *For every $\alpha \in \Sigma$ and every subset U of $T_\alpha \cap Z$ having cardinality less than λ, the following is true: there exist $\beta > \alpha$ and a T_α-T_β linkage $\mathcal{T}$ linking U to B, such that all but fewer than λ paths of $\mathcal{T}$ are contained in paths of $\mathcal{Y}$, and $V(P) \subseteq Z$ for each path $P \in \mathcal{T}$ not contained in a path of $\mathcal{Y}$.*

¿From here the proof continues in a way similar to that of the $\aleph_1$ case. We define inductively ordinals $(\sigma_\alpha : \alpha < \lambda)$, warps $\mathcal{T}_\alpha$ and subsets U_α of T_{σ_α}, as follows. Enumerate $Z \cap A$ as $(z_\alpha : \alpha < \lambda)$ and let $U_0 = \{z_0\}$, $\sigma_0 = 0$. Assume now that σ_α and U_α have been defined. Use Assertion 9.15 to find an ordinal $\beta = \sigma_{\alpha+1} > \sigma_\alpha$ in Σ, and a T_{σ_α}-$T_{\sigma_{\alpha+1}}$-linkage $\mathcal{T}_\alpha$, linking U_α to B and satisfying the conditions stated in the assertion.

Let $U_{\alpha+1}$ consist of the terminal vertex of the path in $*(\mathcal{T}_\theta : \theta \leq \alpha)$ starting at $z_{\alpha+1}$, together with the terminal points of all those paths in $\mathcal{T}_\alpha$ that are not contained in a path of $\mathcal{Y}$.

For limit α let $U_\alpha = ter[*(\mathcal{T}_\theta : \theta < \alpha)\langle\{z_\alpha\}\rangle]$ and $\sigma_\alpha = \sup_{\theta<\alpha} \sigma_\theta$.

Having defined all these for all $\alpha < \lambda$, we define $\mathcal{T} = *(\mathcal{T}_\alpha : \alpha < \lambda)$. For each β, the vertex $z_\beta \in Z \cap A$ is linked to B by $*(\mathcal{T}_\alpha : \alpha \leq \beta)$, and thus it is linked to B by $\mathcal{T}$. Every $a \in A \setminus Z$ is the initial point of some path $P \in \mathcal{F}$. By the way we chose Z, we have $V[\mathcal{T}\langle Z\rangle] \cap V[\mathcal{F}\langle\sim Z\rangle] = \emptyset$ and therefore the set $\mathcal{T}\langle Z\rangle \cup \mathcal{F}\langle\sim Z\rangle$ is a warp. This is the desired A–B-linkage, completing the proof of (♣).

Proof of (♣), Case II: λ is singular.

Definition 9.16. Given a set $\mathcal{P}$ of paths, two vertices u, v are said to be *competitors in* $\mathcal{P}$ if there exist $P, Q \in \mathcal{P}$ such that $in(P) = u$, $in(Q) = v$ and $V(P) \cap V(Q) \neq \emptyset$.

Note that if $\mathcal{P}$ is the union of μ warps, then each vertex has at most μ competitors.

Let $\mathcal{F}$ be a linkage in $(D, A \setminus A', B)$. Let $\mu = cf(\lambda)$ and let $(\kappa_\alpha : \alpha < \mu)$ be a sequence converging to λ. We may assume that $\kappa_0 > \mu$.

Call a matrix of sets *increasing* if each row and each column of the matrix is ascending with respect to the relation of containment.

Assertion 9.17. *There exist two $\mu \times \omega$ matrices: an increasing matrix of sets $(A^k_\alpha : \alpha < \mu, k < \omega)$ and a matrix of half-way linkages $(\mathcal{W}^k_\alpha : \alpha < \mu, k < \omega)$, jointly satisfying the following properties:*

(i) $|A^k_\alpha| = \kappa_\alpha$.
(ii) $\bigcup_{\alpha<\mu} A^0_\alpha = A'$.
(iii) $\mathcal{W}^k_\alpha$ *links* A^k_α *to* B.
(iv) *If* $a \in A^k_\alpha$ *then all competitors of* a *in* $\mathcal{F} \cup \bigcup_{\beta<\mu} \mathcal{W}^k_\beta$ *are in* A^{k+1}_α.
(v) *For every* $\alpha < \mu$ *the sequence* $(\mathcal{W}^k_\alpha : k < \omega)$ *is* $\vec{\preccurlyeq}$*-increasing (as a sequence of warps).*

Proof. We first choose $(A^0_\alpha : \alpha < \mu)$ that satisfy conditions (i) and (ii). We use (♣♣) of the induction hypothesis to obtain half-way linkages $(\mathcal{W}^0_\alpha : \alpha < \mu)$ that satisfy (iii). Denote the stop-over set of $\mathcal{W}^0_\alpha$ by C^0_α. We now define A^1_α to be the set of all competitors of members of A^0_α in $\mathcal{F} \cup \bigcup_{\beta<\mu} \mathcal{W}^0_\beta$. We then use (♣♣) for the webs Γ/C^0_α to get $(\mathcal{W}^1_\alpha : \alpha < \mu)$ that satisfy conditions (iii) and (v). We continue this way, where at each step we define A^{k+1}_α to be the set of all competitors of members of A^k_α in $\mathcal{F} \cup \bigcup_{\beta<\mu} \mathcal{W}^k_\beta$ and we use (♣♣) to get $(\mathcal{W}^{k+1}_\alpha : \alpha < \mu)$ that satisfy conditions (iii) and (v). Condition (i) is satisfied since no vertex has more than μ competitors at any stage.

□

Assertion 9.18. *There exist an ascending sequence of subsets $(A_\alpha : \alpha < \mu)$ of A and a sequence of warps $(\mathcal{W}_\alpha : \alpha < \mu)$, satisfying together the following properties:*

(1) $\mathcal{W}_\alpha$ *links* A_α *to* B.
(2) $\bigcup_{\alpha<\mu} A_\alpha \supseteq A'$.
(3) *If* $a \in A_\alpha$ *then all competitors of* a *in* $\mathcal{F} \cup \bigcup_{\beta<\mu} \mathcal{W}_\beta$ *are also in* A_α.

Proof. Let (A^k_α) and $(\mathcal{W}^k_\alpha)$ be as in Assertion 9.17. Take $A_\alpha = \bigcup_{k<\omega} A^k_\alpha$ and $\mathcal{W}_\alpha = \lim_{k<\omega} \mathcal{W}^k_\alpha$. Conditions (iii) and (v) imply (1), condition (ii) implies (2) and condition (iv) implies (3) because every two competitors in $\mathcal{F} \cup \bigcup_{\beta<\mu} \mathcal{W}_\beta$ are competitors in $\mathcal{F} \cup \bigcup_{\beta<\mu} \mathcal{W}^k_\beta$ for some k.

□

We can now conclude the proof of (♣). For every $a \in \bigcup_{\alpha<\mu} A_\alpha$ use the path to B in $\mathcal{W}_\alpha$ to link a to B, where α is minimal with respect to the property that $a \in A_\alpha$. Such a path exists by condition (1). For every $a \in A \setminus \bigcup_{\alpha<\mu} A_\alpha$, we know by condition (2) that $a \in A \setminus A' = in[\mathcal{F}]$, and hence we can link a to B by the path in $\mathcal{F}$ starting at a. Condition (3) guarantees that these paths are disjoint.

Proof of (♣♣) for general λ (assuming (♣) for cardinals $\leq \lambda$)

Recall that in the case $\lambda = \aleph_0$ and $|V| = \aleph_1$ we used an $\aleph_1$-ladder. Analogously, for general λ we construct a λ^+-ladder, $\mathcal{L}$.

As before, since by Theorem 7.30 $\mathcal{L}$ is not a λ^+-hindrance, there exists a closed unbounded set Σ, disjoint from $\Phi(\mathcal{L})$. Replacing λ by λ^+, we then have the analogues of Corollary 9.8 and Assertions 9.9, 9.12 and 9.14.

The basic idea of the proof is relatively simple. We wish to use (♣) for λ, which is true by the inductive assumption, in order to "climb" $\mathcal{L}$. This is done as follows: Order A' as $(a_i \mid i < \lambda)$. Use Theorem 6.1 to link a_0 to B by a path P so that

$\Gamma - P$ is unhindered. Choose $\alpha_1 \in \Sigma$ such that $V(P) \subseteq RF(T_{\alpha_1})$. Then use Lemma 7.26 and the fact that (♣) holds for λ, to complete P to a linkage $\mathcal{K}_1$ of A into T_{α_1}. Then repeat the procedure with the web Γ_{α_1} replacing Γ, and the element in T_{α_1} to which a_1 is linked by $\mathcal{K}_1$ replacing a_0. After λ such steps, A' is linked to B, and A is linked to some T_γ.

As usual, the problem is the possible generation of infinite paths. To avoid this, we have to anticipate which vertices may participate in infinite paths, and link them to B by the procedure described above. The trouble is that we can take care in this way only of λ^+ such vertices. It is possible for a vertex from A' to have degree larger than λ^+, and then it may be necessary to add more than λ^+ vertices to the set Z of vertices "in jeopardy". The concept used to solve this problem is that of *popularity* of vertices, having in this case a slightly different meaning from the "popularity" of the previous section. "Popularity" of a vertex z means that there exist many z-joined $\mathcal{Y}$-s.a.p's emanating from z, and going to infinity or to B. (In this sense the concept was used in [6] and [9]. A similar notion, solving a similar problem, was used in [5]). A popular vertex does not need to be taken care of immediately, since it can be linked at a later stage, using its popularity. Thus we have to perform the closure operation only with respect to non-popular vertices, and this indeed will necessitate adding only λ^+ vertices to Z.

A first type of vertices which should be considered "popular" are those that do not belong to $RF^\circ(T_\alpha)$ for any $\alpha < \lambda^+$. Note that for each vertex v, the set $\{\theta : \ v \in T_\theta\}$ is an interval, namely it is either empty or of the form $\{\theta : \ \alpha \le \theta < \beta\}$ for some $\alpha < \beta \le \lambda^+$. Let T_{λ^+} be the set of vertices for which this set is unbounded in λ^+. By Lemma 7.16 we have:

Assertion 9.19. *$T_{\lambda^+} = RF(\mathcal{L}) \setminus RF^\circ(\mathcal{L})$.*

As in the proof of (♣) for regular λ, the construction of $\mathcal{L}$ is accompanied by choosing sets Z_α of size at most λ^+, of elements that have to be linked to B in a special way. Let $Z_0 = \emptyset$.

Let $\alpha \le \lambda^+$ (for some definitions below we shall need to refer also to the case $\alpha = \lambda^+$), and assume that we have defined R_β (the rungs of the ladder $\mathcal{L}$) as well as Z_β for all $\beta < \alpha$. Write $Z_{<\alpha} = \bigcup_{\beta<\alpha} Z_\beta$.

Definition 9.20. Let $u \in Z_{<\alpha} \cap RF^\circ(T_\alpha),\ v \in Z_{<\alpha} \cap RF(T_\alpha) \cup \{\infty\}$. A (u, v, α)-*hammock* is a set of pairwise internally disjoint $\mathcal{Y}_\alpha$-s.a.p's from u to v. A (u, v, λ^+)-hammock is plainly called a (u, v)-hammock.

Definition 9.21. Let κ be a cardinality. We say that a (u, v, α)-hammock $\mathcal{H}$ is *maximal up to κ* if one of the following two (mutually exclusive) possibilities occurs:

- $\mathcal{H}$ is a (u, v, α)-hammock which is maximal with respect to inclusion and $|\mathcal{H}| \le \kappa$, or:
- $|\mathcal{H}| = \kappa$ and there exists a (u, v, α)-hammock of size κ^+.

For the construction of Z_α we now choose a (u, v, α)-hammock maximal up to λ^+, for every $u \in Z_{<\alpha} \cap RF^\circ(T_\alpha)$ and every $v \in Z_{<\alpha} \cup \{\infty\}$, and put its entire vertex set into Z_α.

Clearly, a (u, v, α)-hammock that is maximal up to λ^+ contains a (u, v, α)-hammock that is maximal up to μ for every cardinal $\mu < \lambda^+$. Hence, choosing the elements of Z_α carefully, we can see to it that the set $Z = Z_{\lambda^+}$ satisfies:

Assertion 9.22. *For every $\alpha < \lambda^+$, $u \in Z \cap RF(T_\alpha)$, every $v \in (Z \cap RF^\circ(T_\alpha)) \cup \{\infty\}$, and every $\mu \leq \lambda^+$ there exist a (u, v, α)-hammock maximal up to μ, whose vertex set is contained in Z.*

By Theorem 6.1 it is also possible to choose the elements of Z_α so as to guarantee:

Assertion 9.23. *For every $\alpha < \lambda^+$ such that Γ_α is unhindered, and every $v \in T_\alpha \cap Z$, there exists in Γ_α a v-B-path P such that $\Gamma_\alpha - P$ is unhindered and $V(P) \subseteq Z$.*

Yet another condition that can be taken care of is:

Assertion 9.24.

$$V[\mathcal{Y}\langle Z\rangle] \subseteq Z\ .$$

Choosing the vertices y_α of the ladder $\mathcal{L}$ as members of Z, we can ensure:

Assertion 9.25. $Z \subseteq RF(\mathcal{L})$.

Assertion 9.25 will be used to pick objects (like paths or hammocks) contained in Z within $RF(\mathcal{L})$. This will be done without further explicit reference to the assertion.

The description of the construction of $\mathcal{L}$ is now complete. We now show how this construction and the fact that $\Phi = \Phi(\mathcal{L})$ is not stationary can be used to prove the linkability of Γ. As already mentioned, we choose a closed unbounded set Σ disjoint from Φ.

Definition 9.26. A vertex u is said to be *popular* if either $u \in T_{\lambda^+}$, or there exists a (u, ∞)-hammock of cardinality λ^+. The set of popular vertices is denoted by POP.

Remark 9.27. By Lemma 7.17, if $u \in RF(T_\alpha)$, then all $\mathcal{Y}$-alternating paths starting at u are contained in V^α, and are thus $\mathcal{Y}_\alpha$-alternating. Since for each $\alpha < \lambda^+$ we have $|\mathcal{Y}_\alpha\langle\sim A\rangle| \leq \lambda$ and $|\mathcal{Y}_\alpha^\infty| \leq \lambda$, we can assume that all s.a.p's in the hammock witnessing the popularity of u are, in fact, $(\mathcal{Y}_\alpha\langle A\rangle)^f$-alternating.

Let IE be the set of pairs (u, v) of vertices in Z having a (u, v)-hammock of cardinality at least λ^+ ("IE" stands for "imaginary edges"). Let SIE be the set of all pairs (u, v) for which such a hammock exists in which all s.a.p's are non-degenerate (see Definition 4.10), and let $WIE = IE \setminus SIE$ ("SIE" / "WIE" stand for "strong / weak imaginary edges"). Let D' be the graph $(V, E(D) \cup IE)$. Note that possibly $E \cap IE \neq \emptyset$, i.e., there may exist edges that are both "real" and "imaginary".

For a warp $\mathcal{W}$ in D', we define the *real part* $Re(\mathcal{W})$ of $\mathcal{W}$ to be the warp in D whose vertex set is $V[\mathcal{W}]$ and whose edge set is $E[\mathcal{W}] \cap E(D)$. If $u = tail(e)$ for an edge $e \in E[\mathcal{W}] \cap IE$, we write $\mathcal{W}_u$ for the warp obtained from $\mathcal{W}$ by removing e. Also, if $u \in ter[\mathcal{W}]$ we write $\mathcal{W}_u = \mathcal{W}$.

Let us pause to explain the intuition behind these definitions. Consider a warp $\mathcal{W}$ in D' and an imaginary edge $e = (u, v)$ in it. We should think of e as a reminder that we should apply some s.a.p in order to continue the real path ending at u at some later stage of our construction. Since there are λ^+ possible such s.a.p's, not all of them will have been destroyed by the time that it is the turn of u to be linked. Similarly, a popular vertex $v \in ter[\mathcal{W}]$ can wait patiently for its turn to be linked. A vertex $v \in T_{\lambda^+}$ can be linked to B by applying Assertion 9.23 for some α which can be as large as we wish. If there exists a (v, ∞)-hammock of cardinality

λ^+ then, when it is v's turn to be linked, we can use one of the (v,∞)-s.a.p's to link v to T_α for some large $\alpha < \lambda^+$.

Let us now return to the rigorous proof.

Definition 9.28. Given $\alpha \in \Sigma$, a warp $\mathcal{W}$ in D' is called an *α-linkage blueprint* (or α-LB for short) if:

1. $V[\mathcal{W}] \subseteq RF_\Gamma(T_\alpha)$.
2. $in[\mathcal{W} \cup (\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle V[\mathcal{W}]\rangle)] \supseteq A$.
3. $V[\mathcal{W}] \subseteq Z$.
4. $|\mathcal{W}| \leq \lambda$.
5. Every infinite path in $\mathcal{W}$ contains infinitely many strong imaginary edges.
6. $ter[\mathcal{W}] \subseteq POP \cup T_\alpha$.

Definition 9.29. An α-LB $\mathcal{W}$ satisfying $ter[\mathcal{W}] \cap T_\alpha \subseteq T_{\lambda^+}$ is called a *stable* α-LB.

α-linkage blueprints are used to outline a way in which $\mathcal{Y}$ can be altered, via the application of s.a.p's, so as to yield an A-T_α-linkage. An edge $(u,v) \in E[\mathcal{W}] \cap IE$ is going to be replaced by a future application to $\mathcal{Y}$ of a (u,v)-s.a.p. Furthermore, by Definition 9.28(6), terminal vertices of $\mathcal{W}$ not belonging to T_α are popular, again meaning that they can be linked to T_α by the future use of s.a.p's.

Assertion 9.30. *Let $\mathcal{V}$ be an α-LB and let $u \in ter[Re(\mathcal{V})]$. Then there exists an α-LB $\mathcal{G}$ extending $\mathcal{V}_u$, such that $Re(\mathcal{G})$ links u to T_α, and $ter[Re(\mathcal{V})] \subseteq ter[Re(\mathcal{G})] \cup \{u\}$.*

(See Definition 2.3 of a warp being an extension of another warp. Note that in this case, the extension will not necessarily be a forward extension.)

Proof. Let $U = \mathcal{V}(u)$, namely the path in $\mathcal{V}$ containing u. Consider first the case that $u \in ter[\mathcal{V}]$. We may clearly assume that $u \notin T_\alpha$, as otherwise we could take $\mathcal{G} = \mathcal{V}$. By Definition 9.28(6), it follows that $u \in POP$. Since $u \notin T_{\lambda^+}$, by Assertion 9.22 there exists a (u,∞)-hammock $\mathcal{H}$ of size λ^+ contained in Z. Since $|\mathcal{Y}_\alpha\langle\sim A\rangle| \leq \lambda$ and since by Lemma 7.26 also $|\mathcal{Y}_\alpha\langle\sim T_\alpha\rangle| \leq \lambda$, it follows that $\mathcal{H}$ contains a $\mathcal{Y}\langle A, T_\alpha\rangle$-s.a.p Q, that does not meet $V[\mathcal{V}]$ apart from at u. Let $\mathcal{J} = \mathcal{Y} \triangle Q$. Then $\mathcal{G} = \mathcal{V} \diamond \mathcal{J}$ is the desired α-LB (the "$\diamond$" operation is defined in Definition 2.5).

Assume next that $u \notin ter[\mathcal{V}]$. Let (u,v) be the edge in $E[U]$ having u as its tail. Then $(u,v) \in IE$, meaning that there exists a (u,v)-hammock $\mathcal{H}$ of size λ^+, contained in Z. Again, there exists a s.a.p $Q \in \mathcal{H}$ such that $V(Q) \setminus \{u\}$ avoids $\mathcal{Y}_\alpha\langle V[\mathcal{V}]\rangle \cup \mathcal{Y}\langle\sim T_\alpha\rangle$ and $in[\mathcal{J}] \subseteq A$. Let $\mathcal{J} = \mathcal{Y} \triangle Q$. If $(u,v) \in SIE$ we can also assume that $\mathcal{J}$ links u to T_α and hence $\mathcal{V} \diamond \mathcal{J}$ is the desired warp $\mathcal{G}$. If $(u,v) \in WIE$, let $\mathcal{G}_1 = \mathcal{V} \diamond \mathcal{J}$, let P_1 be the path in $Re(\mathcal{G}_1)$ containing u (thus P_1 goes through v, and then continues along U, until it reaches either $ter(U)$ or the next imaginary edge on U), and let $u_1 = ter(P_1)$. Apply the same construction, replacing u by u_1, to obtain an α-LB $\mathcal{G}_2$. By part 5 of definition 9.28 we know that this process will terminate after a finite number of steps. The warp $\mathcal{G}_i$ obtained at that stage is the desired warp $\mathcal{G}$. ☐

We shall need to strengthen Assertion 9.30 in two ways. One is that we wish to link u to B, not merely to T_α. The other is that we wish $\mathcal{G}$ to be a *stable* linkage-blueprint. The next assertion takes care of both these points:

Assertion 9.31. *If $\mathcal{V}$ is an α-LB and $z \in T_\alpha \cap ter[\mathcal{V}]$ then there exist an ordinal $\beta > \alpha$ and a stable β-LB $\mathcal{U}$ extending $\mathcal{V}$, such that:*

(1) *$Re(\mathcal{U})$ links z to B.*
(2) *$ter[Re(\mathcal{V})] \subseteq ter[Re(\mathcal{U})] \cup T_\alpha$.*
(3) *$ter[\mathcal{V}] \cap T_{\lambda^+} \subseteq ter[\mathcal{U}] \cup \{z\}$.*

Proof. By Assertion 9.23 there exists in Γ_α a z-B-path P contained in Z, such that $\Gamma_\alpha - P$ is unhindered.

Claim 1. *There exist a set X of vertices of size at most λ, and an ordinal $\beta > \alpha$, satisfying:*

(1) *$V(P) \cup (ter[\mathcal{V}] \cap T_\alpha) \subseteq X \subseteq Z \cap RF(T_\beta)$.*
(2) *$X \cap T_\beta \subseteq T_{\lambda^+}$.*
(3) *$V[\mathcal{Y}\langle X\rangle] \subseteq X$.*
(4) *$V[\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle T_\beta\rangle] \cup V[\mathcal{Y}\langle T_\beta\rangle \setminus \mathcal{Y}\langle T_\alpha\rangle] \subseteq X$.*
(5) *For every $u \in X \setminus T_{\lambda^+}$ and $v \in X \cup \{\infty\}$ there exists a (u,v)-hammock maximal up to λ contained in X.*

The construction of X and β is done by a closing-up process. By Assertion 9.22, for every $u \in Z \setminus T_{\lambda^+}$ and $v \in Z \cup \infty$ there exists a (u,v)-hammock $H_{u,v}$ contained in Z that is maximal up to λ. Let $M_{u,v} = V[H_{u,v}]$. For $u \in Z \cap T_{\lambda^+}$ let $\gamma_u = \min\{\theta : \ u \in T_\theta\}$. For $u \in Z \setminus T_{\lambda^+}$ define $\gamma_u = \min\{\theta : \ u \in RF^\circ(T_\theta)\}$. For every $\gamma < \lambda^+$ let $H_\gamma = V[\mathcal{Y}\langle T_\alpha\rangle \setminus \mathcal{Y}\langle T_\gamma\rangle] \cup V[\mathcal{Y}\langle T_\gamma\rangle \setminus \mathcal{Y}\langle T_\alpha\rangle]$

Let $\beta_0 = \alpha$ and let $X_0 = V(P) \cup (ter[\mathcal{V}] \cap T_\alpha)$. For every $i < \omega$, let $\beta_{i+1} = \sup\{\gamma_x : \ x \in X_i\}$ and let

$$X_{i+1} = \bigcup_{\substack{u \in X_i \setminus T_{\lambda^+} \\ v \in X_i \cup \{\infty\}}} M_{u,v} \cup H_{\beta_i} \cup V[\mathcal{Y}\langle X_i\rangle] \ .$$

Taking $X = \bigcup_{i<\omega} X_i$ and $\beta = \sup_i \beta_i$ proves the claim.

Claim 2. *Let Q be a (u,v)-s.a.p, where $u \in Z \setminus T_{\lambda^+}$ and $v \in Z \cup \{\infty\}$. If $V(Q) \cap X \subseteq \{u,v\}$ then:*

(1) *If $v \in Z$ then $(u,v) \in IE$.*
(2) *If $v = \infty$ then $u \in POP$.*

To prove (1), assume that $(u,v) \notin IE$. By the properties of X there exists a maximal (u,v)-hammock H lying within X. By the maximality of H, the s.a.p Q must meet some path belonging to H, contradicting the assumption that $V(Q) \cap X = \{u,v\}$. The proof of (2) is similar.

Returning to the proof of the assertion, apply now (♣) to the web $\Gamma_\alpha^\beta - P$, to obtain a T_α-T_β-linkage $\mathcal{W}$ containing P. Let $\mathcal{A} = \mathcal{V} \cup (\mathcal{Y}\langle T_\alpha \cap X, \sim V[\mathcal{V}]\rangle)[RF(T_\alpha)]$ and $\mathcal{C} = \mathcal{A} \diamond \mathcal{W}[X]$. If $V[\mathcal{W}\langle X\rangle] \subseteq X$ then we can take $\mathcal{U} = \mathcal{C}$ to be our desired β-LB. Unfortunately, there is no way to guarantee $V[\mathcal{W}\langle X\rangle] \subseteq X$. Therefore, there might be paths in $\mathcal{W}$ with some vertices in X and some vertices not in X. In this case there may be vertices in $ter[\mathcal{W}[X]]$ which are not in $ter[\mathcal{W}]$, and thus $\mathcal{C}$ might not be a linkage-blueprint, failing to satisfy part 6 of Definition 9.28. This is the reason we need to use imaginary edges. We use imaginary edges to "mend" the holes in $\mathcal{W}[X]$. This is done according to the behavior of $\mathcal{W}$ outside of X.

Define $\mathcal{Z} = \mathcal{W} \downharpoonleft X$, namely the fractured warp consisting of the "holes" formed in $\mathcal{W}$ by the removal of X (thus $E[\mathcal{Z}] = E[\mathcal{W}] \setminus E[\mathcal{W}[X]]$). By Theorem 4.12 and

Remark 4.20 there exists an assignment of an element $v = v(u) \in ter[\mathcal{Z}] \cup \{\infty\}$ and a $(u, v(u))$-$[\mathcal{Z}, \mathcal{Y}]$-s.a.p $Q(u)$ to every $u \in in[\mathcal{Z}]$, such that $v(u_1) \neq v(u_2)$ whenever $u_1 \neq u_2$ and $v(u_1), v(u_2) \in ter[\mathcal{Z}]$.

The desired warp $\mathcal{U}$ is now defined by $ISO(\mathcal{U}) = ISO(\mathcal{V})$ and $E[\mathcal{U}] = E[\mathcal{C}] \cup \{(u, v(u)) \mid u \in in[\mathcal{Z}],\ Q(u) \text{ is finite}\}$. By part (1) of Claim 2 for every u such that $v(u) \in ter[\mathcal{Z}]$ the edge $(u, v(u))$ belongs to IE, and thus $E[\mathcal{U}] \subseteq E \cup IE$. By part (2) of the claim, every $u \in in[\mathcal{Z}]$ for which $v(u) = \infty$ is popular, and thus $ter[\mathcal{U}] \subseteq POP$. By Lemma 4.11, whenever $Q(u)$ is finite and degenerate u and $v(u)$ lie on the same path from $\mathcal{W}$. Since $\mathcal{W}$ is f.c., this implies that every infinite path in $\mathcal{U}$ contains infinitely many non-degenerate edges, as required in the definition of linkage-blueprints. Put together, this shows that $\mathcal{U}$ is a β-LB. By Claim 1(2) it is stable. □

Definition 9.32. For $\alpha \leq \beta < \lambda^+$, we say that a β-LB $\mathcal{U}$ is a *real extension* of an α-LB $\mathcal{V}$ if $Re(\mathcal{U})$ is an extension of $Re(\mathcal{V})$ and $V[\mathcal{V}] \subseteq (ter[\mathcal{U}] \cap ter[\mathcal{V}]) \cup tail[E[\mathcal{U}] \cap E[\mathcal{V}]] \cup V[Re(\mathcal{U})\langle B \rangle]$ We write then $\mathcal{V} \sqsubseteq \mathcal{U}$.

We shall later "grow" blueprints $\mathcal{V}_\alpha$, ordered by the "$\sqsubseteq$" order. The requirement $V[\mathcal{V}] \subseteq (ter[\mathcal{U}] \cap ter[\mathcal{V}]) \cup tail[E[\mathcal{U}] \cap E[\mathcal{V}]] \cup V[Re(\mathcal{U})\langle B \rangle]$ should be thought of as follows. Let $R \in Re(\mathcal{V})$ (so $ter(R)$ is either a vertex in $ter[\mathcal{V}]$ or is the tail of some imaginary edge) and let $R' \in Re(\mathcal{U})$ be the path containing it. One of the following two happens.

- $ter(R) \in ter[Re(\mathcal{U})]$, so $ter(R) = ter(R')$, meaning that R was not "continued forward",
- $ter(R) \in V[Re(\mathcal{U})\langle B \rangle]$, so $ter(R') \in B$, meaning that R was "continued all the way to B".

The third possibility, that R is continued, but not all the way to B, should be disallowed in order to avoid infinite paths.

One can easily check that $\sqsubseteq$ is a partial order. The next assertion states that it behaves well with respect to taking limits:

Assertion 9.33. *Let $\alpha < \lambda^+$ be a limit ordinal and let $(\beta_\theta \mid \theta \leq \alpha)$ be an ascending sequence of ordinals satisfying $\beta_\alpha = \sup_{\theta<\alpha} \beta_\theta < \lambda^+$. Let $\mathcal{V}_\theta$ be a stable β_θ-LB for every $\theta < \alpha$, where $\mathcal{V}_\mu \sqsubseteq \mathcal{V}_\nu$ whenever $\mu < \nu < \alpha$. Let the warp $\mathcal{V}_\alpha = \lim_{\theta<\alpha} \mathcal{V}_\alpha$. Namely, $V[\mathcal{V}_\alpha] = \bigcup_{\theta<\alpha} V[\mathcal{V}_\theta]$ and $E[\mathcal{V}_\alpha] = \bigcup_{\beta<\alpha} \bigcap_{\theta\geq\beta} E[\mathcal{V}_\theta]$ Then $\mathcal{V}_\alpha$ is a stable β_α-LB, that is a real extension of all $\mathcal{V}_{\beta_\theta}$, $\theta < \alpha$.*

Checking most of the properties of an α-LB for $\mathcal{V}_\alpha$ is easy. The only non-trivial part is part (6) of the definition, which follows from the stability of the warps $\mathcal{V}_\theta$.

We can now combine Assertions 9.30 and 9.31, to obtain the following:

Assertion 9.34. *Let $\mathcal{V}$ be a stable α-LB and let $u \in ter[Re(\mathcal{V})]$. Then there exist $\beta > \alpha$ and a stable β-linkage-blueprint $\mathcal{U}$, such that:*

(1) $\mathcal{V} \sqsubseteq \mathcal{U}$.
(2) *$Re(\mathcal{U})$ links u to B, and:*
(3) $ter[Re(\mathcal{V})] \subseteq ter[Re(\mathcal{U})] \cup \{u\}$.

Proof. By Assertion 9.30, there exists an α-LB $\mathcal{G}$ extending $\mathcal{V}_u$, and satisfying $ter[Re(\mathcal{V})] \subseteq ter[Re(\mathcal{G})] \cup \{u\}$. Let z be the terminal vertex of the path in $Re(\mathcal{G})$ containing u. Use Assertion 9.31 to obtain an ordinal $\beta > \alpha$ and a stable β-LB

$\mathcal{U}$ extending $\mathcal{G}$, such that $Re(\mathcal{U})$ links z to B, and $ter[Re(\mathcal{G})] \subseteq ter[Re(\mathcal{U})] \cup T_\alpha$. Thus $ter[Re(\mathcal{V})] \subseteq ter[Re(\mathcal{U})] \cup T_\alpha \cup \{u\}$.

To show that $ter[Re(\mathcal{V})] \subseteq ter[Re(\mathcal{U})] \cup \{u\}$ it suffices to prove that $ter[Re(\mathcal{V})] \cap T_\alpha \subseteq ter[Re(\mathcal{U})] \cup \{u\}$. Note that $ter[Re(\mathcal{V})] \cap T_\alpha \subseteq ter[\mathcal{V}] \cap T_\alpha$. Since $\mathcal{V}$ is a stable α-LB, we have $ter[\mathcal{V}] \cap T_\alpha \subseteq T_{\lambda^+}$. By part (3) of Assertion 9.31, we have $ter[Re(\mathcal{V})] \cap T_\alpha \subseteq ter[Re(\mathcal{U})] \cup \{u\}$. One can easily check the $\mathcal{U}$ is a real extension of $\mathcal{V}$, proving the assertion.

□

We can now conclude the proof of (♣♣). We shall do this by applying Assertion 9.34 λ times. Observe first that $\langle A' \rangle$ is a 0-LB. By Assertion 9.31, it can be extended to a stable σ_0-LB $\mathcal{V}_0$, for some $0 < \sigma_0 < \lambda^+$. Choose now some $u_0 \in ter[Re(\mathcal{V}_0)]$. By Assertion 9.34, there exists a stable σ_1-LB $\mathcal{V}_1$ for some $\sigma_1 > \sigma_0$, such that $\mathcal{V}_0 \sqsubseteq \mathcal{V}_1$ and $Re(\mathcal{V}_1)$ links u_0 to B. We continue this way. For each $\alpha < \lambda$ we choose $u_\alpha \in ter[Re(\mathcal{V}_\alpha)]$ and use Assertion 9.34 to find a stable $\sigma_{\alpha+1}$-LB such that $\mathcal{V}_\alpha \sqsubseteq \mathcal{V}_{\alpha+1}$ and $Re(\mathcal{V}_{\alpha+1})$ links u_α to B. For limit ordinals $\alpha \leq \lambda$ define $\sigma_\alpha = \sup_{\theta<\alpha} \sigma_\theta$ and define $\mathcal{V}_\alpha$ as in Assertion 9.33, so $\mathcal{V}_\alpha$ is a stable σ_α-LB.

Choosing the vertices u_α appropriately, we can procure the following condition:

$$\{u_\alpha : \ \alpha < \lambda\} = \bigcup_{\alpha<\lambda} ter[Re(\mathcal{V}_\alpha)] \setminus B \ .$$

This implies that $\mathcal{V}_\lambda = Re(\mathcal{V}_\lambda)$ and $ter[\mathcal{V}_\lambda] \subseteq B$. Let $\mathcal{H}$ be the warp obtained by adding to $\mathcal{V}_\lambda$ all paths of $\mathcal{Y}$ not intersecting $V[\mathcal{V}_\lambda]$ and let $\sigma = \sigma_\lambda$. Then $\mathcal{H}$ is an A-T_σ-linkage linking A' to B. Since Γ/T_σ is unhindered, $\mathcal{H}$ is a half-way linkage, as required in the theorem. □

10. Open problems in infinite matching theory

The Erdős-Menger conjecture pointed at the way duality should be formulated in the infinite case: rather than state equality of cardinalities, the conjecture stated the existence of dual objects satisfying the so-called "complementary slackness conditions". There are still many problems of this type that are open. One of the most attractive of those is the "fish-scale conjecture", named so because of the way its objects can be drawn [10]:

Conjecture 10.1. *In every poset not containing an infinite antichain there exist a chain C and a decomposition of the vertex set into antichains A_i, such that C meets every antichain A_i.*

The dual statement, obtained by replacing the terms "chain" and "antichain", follows from the infinite version of König's theorem [26, 7]. It is likely that, if true, Conjecture 10.1 does not have much to do with posets, but with a very general property of infinite hypergraphs.

Definition 10.2. Let $H = (V, E)$ be a hypergraph. A *matching* in H is a subset of E consisting of disjoint edges. An *edge cover* is a subset of E whose union is V. A matching I is called *strongly maximal* if $|J \setminus I| \leq |I \setminus J|$ for every matching J in H. An edge cover F is called *strongly minimal* if $|K \setminus F| \geq |F \setminus K|$ for every edge cover K in H.

As noted above, our main theorem is tantamount to the fact that the hypergraph of vertex sets of A–B-paths in a web possesses a strongly maximal matching. Call a hypergraph *finitely bounded* if its edges are of size bounded by some fixed finite number. Call a hypergraph H a *flag complex* if it is closed down, namely every subset of an edge is also an edge, and it is 2-determined, namely if all 2-subsets of a set belong to H then the set belongs to H.

Conjecture 10.3.

(1) *Every finitely bounded hypergraph contains a strongly maximal matching and a strongly minimal cover.*
(2) *Any flag complex contains a strongly minimal cover.*

Conjecture 10.1 would follow by a compactness argument from part (2) of this conjecture. For graphs part (1) of the conjecture follows from the main theorem of [5].

The mere condition of having only finite edges does not suffice for the existence of a strongly maximal matching, as was shown in [12]. In the example given there, for every matching M there exists a matching M' with $|M \setminus M'| = 2,\ |M' \setminus M| = 3$.

Problem 10.4 (Tardos)**.** *Is it true that in every hypergraph with finite edges there exists a matching M such that no matching M' exists for which $|M \setminus M'| = 1,\ |M' \setminus M| = 2$?*

Acknowledgement We are grateful to the members of the Hamburg University Combinatorics seminar led by Reinhard Diestel, for a careful reading of a preliminary draft of this paper, and for pointing out many inaccuracies. In particular, Henning Bruhn and Maya Stein contributed to the presentation and correctness of Sections 8 and 9, and Philipp Sprüssel made helpful remarks concerning Section 2. We are also grateful to Eliahu Levy and Maté Vizer for many useful comments.

References

[1] R. Aharoni, Menger's theorem for graphs containing no infinite paths, *Eur. J. Combin.*, **4**(1983), 201–204.
[2] R. Aharoni, König's duality theorem for infinite bipartite graphs, *J. London Math. Soc.*, **29** (1984),1–12.
[3] R. Aharoni, Menger's theorem for countable graphs. *J. Combin. Th., Ser. B*, **43**(1987), 303–313.
[4] R. Aharoni, Matchings in graphs of size $\aleph_1$, J. Combin. Theory Ser. B **36** (1984), 113–117.
[5] R. Aharoni, Matchings in infinite graphs, *J. Combin. Th.* Ser. B, **44**(1988), 87–125.
[6] R. Aharoni, Linkability in countable-like webs, *Cycles and Rays, G. Hahn et al. eds., Nato ASI Series*, 1990, 1–8.
[7] R. Aharoni, Infinite matching theory, *Disc. Math.* **95**(1991), 5–22.
[8] R. Aharoni, A few remarks on a conjecture of Erdős on the infinite version of Menger's theorem, *The Mathematics of Paul Erdős*, R.L. Graham and J. Nesteril eds., Springer-Verlag Berlin Heidelberg, 1997, 394–408.
[9] R. Aharoni, and R. Diestel, Menger's theorem for countable source sets, *Comb., Probability and Comp.* **3**(1994), 145–156.
[10] R. Aharoni and V. Korman, Greene-Kleitman's theorem for infinite posets, *Order* **9** (1992), 245–253.
[11] R. Aharoni, C.St.J.A Nash-Williams and S. Shelah, A general criterion for the existence of transversals, *Proc. London Math. Soc.* **47**(1983), 43–68.
[12] R. Ahlswede and L. H. Khachatrian, A counterexample to Aharoni's strongly maximal matching conjecture, *Discrete Math.* **149**(1996), 289.
[13] M.R. Damerell and E.C. Milner, Necessary and sufficient conditions for transversals of countable set systems, *J. Combin. Th. Ser. A* **17**(1974), 350–374
[14] R. Diestel, Graph Theory, *Springer-Verlag* (1997, 1-st edition).
[15] T. Gallai, Ein neuer Bewies eines Mengerschen Satzes, *J. London Math. SOc.* **13**(1938)

[16] P. Hall, On representatives of subsets, *J. London Math. Soc.* **10**(1935), 26-30.
[17] D. König, Graphs and matrices, *Mat. Fiz. Lapok* **38**, 1931, 116-119. (Hungarian)
[18] D. König, *Theorie der endlichen und unendlichen Graphen*, Akademischen Verlagsgesellschaft, Leipzig, 1936. (Reprinted: Chelsea, New York, 1950.)
[19] L. Lovász and M. D. Plummer, *Matching Theory*, Annals of Mathematics 29, North Holland, 1991.
[20] C. McDiarmid, On separated separating sets and Menger's theorem. *Congressus Numerantium* **15**(1976), 455-459.
[21] K. Menger, Zur allgemeinen Kurventhoerie, *Fund. Math.* **10**(1927), 96–115.
[22] C. St.J. A. Nash-Williams, Infinite graphs—a survey. *J. Combin. Th.* **3** (1967), 286–301.
[23] C. St.J. A. Nash-Williams, Which infinite set-systems have transversals?—a possible approach. Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972), 237–253. Inst. Math. Appl., Southend-on-Sea, 1972.
[24] C. St.J. A. Nash-Williams, Marriage in denumerable societies, *J. Combin. Th.* Ser. A **19**(1975), 335–366.
[25] C. St.J. A. Nash-Williams, Another criterion for marriage in denumerable societies, in: B. Bollobás, ed., Advances in Graph Theory, *Ann. Discrete Math.* **3**, North-Holland, Amsterdam, 1978, 165–179.
[26] H. Oellrich and K. Steffens, On Dilworth's decomposition theorem, *Discrete Math.* **15**(1976), 301–304.
[27] K.-P. Podewski and K. Steffens, Injective choice functions for countable families. *J. Combin. Th. Ser. B* **21** (1976), 40–46.
[28] K.-P. Podewski and K. Steffens, Uber Translationen und der Satz von Menger in unendlichen Graphen, *Acta Math. Sci. Hungar.* **30**(1977), 69-84.
[29] Theory of Graphs and its Applications, Proceedings of the Symposium held in Smolenice, June 1963, Czechoslovak Academy of Sciences, Prague, 1964.

Department of Mathematics, Technion, Haifa, Israel 32000
E-mail address, Ron Aharoni: `ra@tx.technion.ac.il`

Department of Mathematics, Haifa University, Haifa, Israel
E-mail address, Eli Berger: `berger@math.haifa.ac.il`