HAMBURGER BEITRÄGE
ZUR MATHEMATIK

Heft 291

The Max-Flow Min-Cut Theorem for Countable Networks

Ron Aharoni
Eli Berger
Agelos Georgakopoulos
Amitai Perlstein
Philipp Sprüssel

Version November 2007

THE MAX-FLOW MIN-CUT THEOREM FOR COUNTABLE NETWORKS

RON AHARONI, ELI BERGER, ANGELOS GEORGAKOPOULOS, AMITAI PERLSTEIN, AND PHILIPP SPRÜSSEL

Abstract. We prove a strong version of the the Max-Flow Min-Cut theorem for countable networks, namely that in every such network there exist a flow and a cut that are "orthogonal" to each other, in the sense that the flow saturates the cut and is zero on the reverse cut. If the network does not contain infinite trails then this flow can be chosen to be *mundane*, i.e. to be a sum of flows along finite paths. We show that in the presence of infinite trails there may be no orthogonal pair of a cut and a mundane flow. We finally show that for locally finite networks there is an orthogonal pair of a cut and a flow that satisfies Kirchhoff's first law also for ends.

1. Introduction

Recently, the first two authors of this paper proved the following generalization of Menger's theorem to the infinite case [5]:

Theorem 1.1. *Given a possibly infinite digraph and two vertex sets A and B in it, there exists a set P of vertex-disjoint A-B paths and an A-B-separating set of vertices S, such that S consists of a choice of precisely one vertex from each path in P.*

In the finite case, the closely related edge version of Menger's theorem can be viewed as the integral version of the Max-Flow Min-Cut (MFMC) theorem. In fact, the MFMC theorem can easily be reduced to Menger's theorem, while the standard proofs of the MFMC theorem yield also its integral version, namely the edge version of Menger's theorem.

Thus it is natural to ask also for a generalization of the MFMC theorem to the infinite case. Theorem 1.1, which was originally conjectured by Erdős, suggests a possible generalization. In the language of Linear Programming, the infinite version of Menger's theorem is formulated in terms of the complementary slackness conditions, rather than of equality of the values of dual programs. Applying this to the MFMC theorem we are naturally led to the conjecture that in any network there exists an *orthogonal pair* of a flow and a cut, i.e. a flow and a cut related to each other by the complementary slackness conditions; these demand that every edge of the cut is saturated by the flow, and on each edge of the reverse cut the value of the flow is zero (see Section 2 for precise definitions).

The research of the first author was supported by grants from the Israel Science Foundation, BSF, the M. & M.L Bank Mathematics Research Fund and the fund for the promotion of research at the Technion and a Seniel Ostrow Research Fund.

The research of the second author was supported by a BSF grant.

The research of the first, third and fifth authors was supported by a GIF grant.

Conjecture 1.2. *In any (possibly infinite) network there exists an orthogonal pair of a flow and a cut.*

In this paper we prove this conjecture for countable networks (Section 6). For this proof we will use a generalization of the notion of waves, which was a central tool in the proof of Theorem 1.1.

This conjecture however does not generalize Theorem 1.1 even in its integral version, since the flow may contain infinite paths. This naturally raises the question whether Conjecture 1.2 is true for flows that do not allow any flow to escape to infinity; we will call such flows *mundane* (see Section 7 for the precise definition). As we shall see, this is not the case, since even in locally finite networks mundane flows do not necessarily attain the supremum of their values. We thus seek to relax the constraint that no flow escapes to infinity, and are led to two new types of flows: *finite-cut-respecting* flows, which are allowed to send flow to infinity but any amount flowing into an end of the graph must flow out of the same end, and *cut-respecting* flows, which are finite-cut-respecting flows with the additional constraint that if flow circumvents some cut F by flowing through an end, then this circumvention does not exceed the amount that could in principle flow through F (see Section 8 for precise definitions).

For each of these kinds of flows we show, at least for locally finite networks, that the infimum of the capacities of all relevant cuts equals the supremum of the values of the corresponding flows (Sections 7 and 8). Moreover, we decide when these infima and suprema are attained by some cut or flow (Sections 5, 7 and 8). Finally, we prove that in any locally finite network there is an orthogonal pair of a cut-respecting flow and a cut of minimum capacity (Section 8).

2. Definitions and Notation

As a general rule, we shall use the terminology of [1]. Deviations will be explicitly indicated.

The characteristic function of a subset T of a set S is denoted by $\chi_S(T)$, or simply $\chi(T)$ if the identity of S is clear from the context.

For a directed edge $e = (u, v)$ we shall write $u = init(e)$, $v = ter(e)$. For a vertex v in a digraph we denote by $OUT(v)$ the set of edges e with $init(e) = v$, and by $IN(v)$ the set of edges e with $ter(e) = v$.

Definition 2.1. A *network* Δ is a quadruple (D, c, s, t), where $D = (V, E)$ is a digraph with no loops, c is a function (called *capacity*) from E to $\mathbb{R}_+$, and s, t are vertices of D, called *source* and *sink* respectively. We shall assume that $IN(s) = OUT(t) = \emptyset$.

Let $\Delta = (D, c, s, t)$ where $D = (V, E)$ be a network fixed throughout this section.

For a function g on E and an edge $(u, v) \in E$ we abbreviate $g((u, v))$ to $g(u, v)$. If F is a set of edges we write $g[F] = \sum\{g(e) \mid e \in F\}$. We adopt the convention that if $\sum\{g(e) \mid e \in F,\ g(e) > 0\} = \infty$ and $\sum\{g(e) \mid e \in F,\ g(e) < 0\} = -\infty$ then $g[F] = 0$.

For a vertex $v \in V$ we write $d_g^-(v) = g[IN(v)]$, $d_g^+(v) = g[OUT(v)]$, and $d_g(v) = d_g^+(v) - d_g^-(v)$. Here, again, we adopt the convention $\infty - \infty = 0$. Given a function f on the edge set of an undirected graph, the *degree* $d_f(v)$ of a vertex v is the sum of $f(e)$ over all edges e incident with v.

Definition 2.2. Given a function f on E, the set of vertices $x \in V$ for which $d^+_f(x) = 0$ holds is denoted by $SNK(f)$. The set of vertices x for which $d_f(x) = 0$ (and thus $d^+_f(x) = d^-_f(x)$) holds is denoted by $KIR(f)$ (KIR standing for "Kirchhoff").

Definition 2.3. A function $f : E \to \mathbb{R}_+$ is called a *flow* if:

- (Capacity constraint:) $f(e) \leq c(e)$ for every $e \in E$.
- (Flow conservation:) $V \setminus \{s, t\} \subseteq KIR(f)$.

The set of edges e for which $f(e) > 0$ holds is denoted by *supp*(f). The *value* $|f|$ of a flow f is defined by $|f| := d^+_f(s)$. Note that in infinite networks this is not necessarily equal to $d^-_f(t)$. If $\vec{C}$ is a directed in D then we say that $\vec{C}$ is a *cycle of* f if $f(e) > 0$ holds for every edge $e \in E(\vec{C})$.

A *cut* is a set of edges of the form $E(S, V \setminus S)$ for some $S \subseteq V$, where $E(X, Y)$ is the set of edges directed from X to Y. An s–t cut is a cut $E(S, V \setminus S)$ such that $s \in S$ and $t \notin S$. A flow f is said to *saturate* an edge e if $f(e) = c(e)$. It is said to saturate a set F of edges if it saturates all edges in F. A flow f and an s–t cut $E(S, V \setminus S)$ are *orthogonal* to each other if f saturates $E(S, V \setminus S)$ and is zero on every edge in $E(V \setminus S, S)$.

A 1-way infinite path in a graph G is called a *ray*. Two rays R, L in G are *equivalent* if no finite set of vertices separates them. The corresponding equivalence classes of rays are the *ends* of G.

3. A VERTEX VERSION

As already mentioned, in the finite case the edge version of Menger's theorem is just the integral case of the MFMC theorem, namely the case in which the capacity function is identically 1 and the desired flow only takes the values 0 and 1. The vertex version and the edge version of Menger's theorem are easily derivable from each other. To get the vertex version from the edge version, one splits each vertex into a "receiving" copy and an "emitting" copy, connected by an edge. The derivation of the edge version from the vertex version is done by a transformation that can be used also in the non-integral case, yielding an equivalent version of Conjecture 1.2. To state it we need the following definitions:

Definition 3.1. A *weighted web* Γ is a quadruple (D, A, B, w), where D is a digraph, $A, B \subseteq V(D)$ and w is a function from $V(D)$ to $\mathbb{R}_+$. Let $V(\Gamma) = V(D)$ and $E(\Gamma) = E(D)$.

Let $\Gamma = (D, A, B, w)$ be a weighted web fixed throughout this section.

Definition 3.2. A *current* f in Γ is a function from $E(D)$ to $\mathbb{R}_+$ such that

(i) $d^+_f(x) \leq w(x)$ and $d^-_f(x) \leq w(x)$ for every vertex $x \in V(D)$;
(ii) $d^+_f(x) \leq d^-_f(x)$ for every vertex $x \in V(D) \setminus A$; and
(iii) $d^-_f(a) = 0$ for every $a \in A$ and $d^+_f(b) = 0$ for every $b \in B$.

A vertex x is said to be *saturated* by f if $x \in A$ or $d^-_f(x) = w(x)$. The set of vertices that are saturated by f is denoted by $SAT(f)$. The set $SAT(f) \cap SNK(f)$ is denoted by $TER(f)$ (standing for "terminal points"; recall that $SNK(f)$ is the set of vertices x for which $d^+_f(x) = 0$).

Definition 3.3. A current f satisfying $KIR(f) \supseteq V(D) \setminus (A \cup B)$ is called a $web - flow$.

A set S of vertices in Γ is said to be *A-B-separating* (or simply *separating*) if every path from A to B meets S. Given a (not necessarily separating) subset S of $V(D)$, a vertex $x \in S$ is said to be *essential (for separation)* in S if it is not separated from B by $S \setminus \{x\}$. The set of essential elements of S is denoted by $\mathcal{E}(S)$. It is easy to show:

Lemma 3.4 ([5])**.** *If S is separating, then so is $\mathcal{E}(S)$.*

For a set S of vertices in Γ we denote by $RF(S) = RF_\Gamma(S)$ the set of vertices separated by S from B. (The letters "RF" stand for "roofed", a term originating in the way the authors draw their weighted webs, with the "A" side at the bottom, and the "B" side on top.) In particular, $S \subseteq RF(S)$. We define $RF^\circ(S) := RF(S) \setminus \mathcal{E}(S)$.

Given a current f, we write $RF(f) = RF(TER(f))$ and $RF^\circ(f) = RF^\circ(TER(f))$.

Definition 3.5. Let f be a web flow and let S be a separating set. We say that S is *orthogonal* to f if $S \subseteq SAT(f)$ and $f(u,v) = 0$ for every pair of vertices u, v with $v \in RF(S)$ and $u \in V \setminus RF^\circ(S)$.

An equivalent conjecture to Conjecture 1.2 is:

Conjecture 3.6. *In every weighted web there exists a web-flow f and an A-B separating set orthogonal to f.*

The transformation used to deduce Conjecture 1.2 from Conjecture 3.6 is the following. Let Δ be a network, with notation as in Definition 2.1. Let $\Gamma = (D', A, B, w)$ be the web defined by $V(\Gamma) = E(\Delta)$, $E(\Gamma) = \{((x,y),(y,z)) \mid (x,y),(y,z) \in E(\Delta)\}$, $A = E(s, V(\Delta) \setminus \{s\})$, $B = E(V(\Delta) \setminus \{t\}, t)$, and $w(e) = c(e)$ for every $e \in V(\Gamma) = E(\Delta)$.

Easily, every A-B-separating set of vertices in Γ is also an s-t-separating set of edges in Δ. If f is a web-flow in Γ, we can define a flow g in Δ as $g(e) = max(d_f^+(e), d_f^-(e))$ (recall, however, that $d_f^+(e) = d_f^-(e)$ if $e \notin A \cup B$). It is straightforward to check that g is indeed a flow. Moreover, if f is orthogonal to some A-B-separating set of vertices then g is orthogonal to the corresponding set of edges in Δ.

In the following sections we will prove Conjecture 3.6, and thus Conjecture 1.2, for the countable case (see Theorem 6.1).

4. Linkability in weighted webs, waves, and an equivalent conjecture

In this section we develop some tools that we will use for the proof of Conjecture 3.6 for countable weighted webs. These are generalizations of fundamental notions in the proof of Theorem 1.1 in [5], and they could turn out useful in proving the general case of Conjecture 3.6.

Let $\Gamma = (D, A, B, w)$ be a weighted web fixed throughout this section.

Definition 4.1. A web-flow f in Γ is called a *linkage* if $d_f^+(a) = w(a)$ for every $a \in A$. If a weighted web contains a linkage it is called *linkable*.

Definition 4.2. A current f in Γ is called a *wave* if $TER(f)$ is A-B-separating and $d_f^+(x) = 0$ for all $x \notin RF(f)$.

If f, g are waves, we write $f \leq g$ if $f(e) \leq g(e)$ for every edge e.

Lemma 4.3. *Let I be a totally ordered set, and let $(f_i \mid i \in I)$ be waves such that $f_i \leq f_j$ whenever $i \leq j$. Then $f = \sup(f_i \mid i \in I)$ is a wave.*

Proof. Let P be an A-B path. Easily, for $j \geq i$ we have $SNK(f_i) \supseteq SNK(f_j)$ and $SAT(f_i) \subseteq SAT(f_j)$. As P is finite, this means that there exists an i such that for every $j \geq i$ we have $SNK(f_i) \cap V(P) = SNK(f_j) \cap V(P)$ and $SAT(f_i) \cap V(P) = SAT(f_j) \cap V(P)$. Then, $SNK(f) \cap V(P) = SNK(f_i) \cap V(P)$ and $SAT(f) \cap V(P) \supseteq SAT(f_i) \cap V(P)$. Hence, $TER(f) \cap V(P) \supseteq TER(f_i) \cap V(P)$, and since f_i is a wave, this implies that $TER(f) \cap V(P) \neq \emptyset$. This proves that f is a wave. $\square$

By Zorn's lemma this implies:

Corollary 4.4. *In every weighted web there exist a $(\leq)$-maximal wave.*

Definition 4.5. A wave f is called a *hindrance* if there exists a vertex $a \in A \setminus \mathcal{E}(TER(f))$ such that $d^+_f(a) < w(a)$. If $0 < \varepsilon < w(a) - d^+_f(a)$ then f is said to be a $(> \varepsilon)$-hindrance. A weighted web is called *hindered* (respectively *$(> \varepsilon)$-hindered*) if it contains a hindrance (respectively a $(> \varepsilon)$-hindrance). A weighted web is called *loose* if it contains no non-zero wave and the zero wave is not a hindrance.

The following is an easy consequence of the definitions.

Observation 4.6. *Let $\Gamma = (D, A, B, w)$ and $\Gamma' = (D, A, B, w')$ be weighted webs such that $w'(v) = w(v)$ for all $v \in V \setminus A$ and $w'(a) \leq w(a)$ for all $a \in A$. Then every wave in Γ' is a wave in Γ. Thus if Γ is loose then so is Γ'.*

Definition 4.7. Let f be a wave. A wave g is called a *trimming* of f if

(i) $g \leq f$
(ii) $RF^\circ(f) \subseteq KIR(g) \cup A$ and:
(iii) $TER(g) \setminus A = \mathcal{E}(TER(f)) \setminus A$.

A wave is called *trimmed* if it is a trimming of itself.

Lemma 4.8. *Every wave has a trimming.*

Proof. Let f be a wave that is not trimmed, let $x \in RF^\circ(f) \setminus (KIR(f) \cup A)$ and let f_1 be the wave obtained from f by decreasing the values on $IN(x)$ so that $d^-_f(x) = d^+_f(x)$. One can easily see that $\mathcal{E}(TER(f_1)) = \mathcal{E}(TER(f))$, which means that f_1 is indeed a wave. If f_1 is trimmed, we are done. If not, we can find in a similar way a wave $f_2 \leq f_1$ with $\mathcal{E}(TER(f_2)) = \mathcal{E}(TER(f))$. Continuing this process, if necessary, transfinitely, we obtain a trimmed wave f_α, which is then a trimming of f. $\square$

Definition 4.9. If Γ is a weighted web and f is a wave in Γ, we write Γ/f for the web Ξ defined by $A_\Xi = \mathcal{E}(TER(f))$, $B_\Xi = B$, $V_\Xi = V \setminus RF^\circ(f)$, $D_\Xi = D[V_\Xi]$ (the subgraph of D induced on V_Ξ) and $w_\Xi = w \restriction V_\Xi$.

Waves can be combined, as follows:

Definition 4.10. Let f be a wave and g be a current. We denote by $f^\frown g$ the function $f + (g \restriction E(\Gamma/f))$.

It is easy to check that $f^\frown g$ is a current. In fact, if g is a wave, then $g \restriction (\Gamma/f)$ is a wave in Γ/f, and thus $f^\frown g$ is a wave, which follows from:

Lemma 4.11. *If $g \restriction (\Gamma/f)$ is a wave, then $TER(f \frown g) \supseteq \mathcal{E}(TER(f) \cup TER(g))$.*

Proof. Let $x \in \mathcal{E}(TER(f) \cup TER(g))$. We wish to show that $x \in TER(f \frown g)$.

We first note that x cannot lie in $RF^\circ(f)$ or $RF^\circ(g)$, so $x \in SNK(f) \cap SNK(g)$ hence $x \in SNK(f \frown g)$. It remains to show that $x \in SAT(f \frown g)$. Since $f \frown g \geq f$ we have $SAT(f \frown g) \supseteq SAT(f)$. Hence we are done in the case $x \in TER(f)$ and we may assume $x \in TER(g) \setminus TER(f)$.

Since $x \notin TER(f)$ and $x \notin RF^\circ(f)$, we have $x \notin RF(f)$ and thus, for an edge e entering x we have $e \in E(\Gamma/f)$ and thus $(f \frown g)(e) = g(e)$. Since $x \in SAT(g)$ this yields $x \in SAT(f \frown g)$, completing the proof. □

The following lemma is easy to prove:

Lemma 4.12. *If f and g are waves in a bipartite web, then $TER(f \frown g) \cap B = (TER(f) \cup TER(g)) \cap B$.*

We can now use our new machinery to reformulate Conjecture 3.6:

Conjecture 4.13. *A loose weighted web is linkable.*

Lemma 4.14. *Conjecture 4.13 implies Conjecture 3.6 and hence Conjecture 1.2.*

Proof. Let Γ be a weighted web, and let f be a ($\leq$)-maximal wave in Γ. Let $T = \mathcal{E}(TER(f))$, and let h be a trimming of f. Clearly, Γ/f is loose. Assuming Conjecture 4.13, there exists a linkage g in Γ/f. Then, $k = h + g$ is a web-flow, $T \subseteq SAT(k)$ and T is A-B separating. Since $supp(g) \subseteq V(\Gamma/f)$ we have $k(x, y) = 0$ for every pair of vertices x, y with $x \in V \setminus RF^\circ(T)$ and $y \in RF(T)$, which proves that T is orthogonal to k. □

5. Attainability of flow values in infinite networks

In this section we return to flows in networks, rather than web-flows. Our aim is to prove a result which will serve as a main ingredient in the proof of Conjecture 3.6 for countable weighted webs (Theorem 6.1), and which seems to be of independent interest:

Theorem 5.1. *In a countable network Δ where $d^-_c(x) < \infty$ (i.e. the sum of the capacities of the edges pointing to x is finite) for every vertex x, there exists a flow f such that $|f| = sup\{|g| : g \text{ is a flow in } \Delta\}$ and $d^-_f(x) \leq |f|$ for every vertex x. In particular, if the values of flows in Δ are unbounded, then there exists a flow of infinite value.*

Definition 5.2. Let f be a flow in a network $\Delta = (D, c, s, t)$ that contains no pair of edges with the same endvertices but opposite directions. The *residual network* $RES(\Delta, f)$ of Δ and f is the network (D', c_R, s, t) where D' is the digraph obtained from D by adding an edge (u, v) for every edge $(v, u) \in E(D) \setminus (OUT(s) \cup IN(t))$, and where c_R is defined by letting, for every edge $(x, y) \in E(D)$, $c_R(x, y) := c(x, y) - f(x, y)$ and $c_R(y, x) := f(x, y)$.

For a function g on the edge-set of $RES(\Delta, f)$ let $f \oplus g$ denote the function h on the edge-set of Δ defined by $h(x, y) = f(x, y) + g(x, y) - g(y, x)$. The following is a straightforward corollary of the definitions:

Lemma 5.3. *Let f be a flow in Δ and let g be a flow in $RES(\Delta, f)$. Then $f \oplus g$ is a flow in Δ, with $|f \oplus g| = |f| + |g|$.*

The following result shows that it is possible to clean up a flow from cycles and current coming from infinity without reducing its value.

Lemma 5.4. *If g is a flow in Δ of finite value then there exists a flow $h \leq g$ such that $|h| = |g|$ and $d^-_h(x) \leq |h|$ for every vertex x.*

Proof. We propagate the desired flow h from s. We will define h recursively in infinitely many steps, in each step considering a vertex and adjusting its out-degree to its in-degree, and then removing any cycles we created in doing so. But as subsequent steps might change the in-degree of a vertex we already considered, we will have to return to each vertex infinitely often.

Formally, let $v_1, v_2, \ldots$ be a sequence in which each vertex in $V(\Delta)\backslash\{s,t\}$ appears infinitely often. We will recursively define sequences (h_i), (h^+_i) and (h^-_i) of functions on $E(\Delta)$. Intuitively, h^+_i differs from h^+_{i-1} in that it makes the out-degree of v_i equal to its in-degree, while h^-_i is the sum of some unwanted cycles in h^+_i. Subtracting the two we obtain the functions h_i that will converge to the desired flow h.

While defining these sequences we will make sure that the following conditions are satisfied for all $i \in \mathbb{N}$:

(i) If $i > 0$, then $h^+_{i-1}(e) \leq h^+_i(e)$ and $h^-_{i-1}(e) \leq h^-_i(e)$ for every edge e;
(ii) $h^-_i(e) \leq h^+_i(e) \leq g(e)$ for every edge e;
(iii) the support of $h_i := h^+_i - h^-_i$ does not contain cycles;
(iv) $d^+_{h^+_i}(v) \leq d^-_{h^+_i}(v)$ for every $v \in V(\Delta) \setminus \{s,t\}$; and
(v) $d^+_{h^-_i}(v) = d^-_{h^-_i}(v)$ for every $v \in V(\Delta) \setminus \{s,t\}$.

We start by defining $h^+_0 = g$ on $OUT(s)$ and $h^+_0 = 0$ on all other edges, and $h^-_0 = 0$. Clearly, conditions (i)–(v) are satisfied for $i = 0$. For $i = 1, 2, \ldots$, assume that h^+_j and h^-_j have already been defined for every $j < i$ and satisfy (i)–(v).

We define h^+_i first. If $d^+_{h^+_{i-1}}(v_i) < d^-_{h^+_{i-1}}(v_i)$, then give each edge e in $OUT(v_i)$ a value $h^+_i(e)$ with $h^+_{i-1}(e) \leq h^+_i(e) \leq g(e)$ so that (having considered all edges in $OUT(v_i)$) $d^+_{h^+_i}(v_i) = d^-_{h^+_{i-1}}(v_i)$ holds; this is possible since $d^+_g(v_i) = d^-_g(v_i)$ and $h^+_{i-1}(e) \leq g(e)$ for every $e \in E(\Delta)$ by condition (ii). For every edge e in $E(\Delta) \setminus OUT(v_i)$ let $h^+_i(e) = h^+_{i-1}(e)$. Clearly, conditions (i)–(v) are not violated.

Next we define h^-_i. The function $h'_i := h^+_i - h^-_{i-1}$ is non-negative since $h^+_i \geq h^+_{i-1} \geq h^-_{i-1}$ by (i) and (ii). If $supp(h'_i)$ contains any cycles, then let $C_1, C_2, \ldots$ be a (possibly infinite) enumeration of those cycles. (The C_j are not necessarily pairwise edge disjoint.) We are going to remove all cycles from $supp(h'_i)$ by performing infinitely many steps (within step i), in each step j eliminating the cycle C_j from $supp(h'_i)$. For every edge e we denote by $h^j_i(e)$ the value that has to be subtracted from $h'_i(e)$ in order to eliminate the cycles $C_1, \ldots, C_j$. To begin with, let $h^0_i(e) = 0$ for every $e \in E(\Delta)$. For $j = 1, 2, \ldots$, if C_j is a cycle in $supp(h'_i - h^{j-1}_i)$ then, for every edge $e \in E(C_j)$, add to $h^{j-1}_i(e)$ the value

$$\min\{h'_i(e) - h^{j-1}_i(e) \mid e \in E(C_j)\}$$

to obtain $h^j_i(e)$; let $h^j_i(d) = h^{j-1}_i(d)$ for every other edge d.

Having treated all cycles C_j, define $h^-_i(e) := h^-_{i-1}(e) + lim_j h^j_i(e)$ for every $e \in E(\Delta)$; this is well defined since $h^j_i(e)$ is monotone increasing with j and bounded by $h'_i(e)$. It is not hard to see that conditions (i)–(v) are satisfied.

By (i) and (ii) the sequences $h_i^+(e)$, $h_i^-(e)$ and $h_i(e)$ converge for every edge e; we let $h^+(e) := \lim_i h_i^+(e)$, $h^-(e) := \lim_i h_i^-(e)$, and $h(e) := \lim_i h_i(e) = h^+(e) - h^-(e)$. By (ii) we have $h^+(e) \leq g(e)$. Since for every vertex $v \in V(\Delta) \setminus \{s,t\}$ we have $d^+_{h_i^+}(v) = d^-_{h_i^+}(v)$ for infinitely many i, we have $d^+_{h^+}(v) = d^-_{h^+}(v) \leq d^-_g(v) < \infty$. By (i), (ii) and (v) we have $d^+_{h^-}(v) = d^-_{h^-}(v) \leq d^-_{h^+}(v)$. Hence h is non-negative and $d^+_h(v) = d^-_h(v)$, and therefore h is indeed a flow.

As $IN(s) = \emptyset$, for every edge $e \in OUT(s)$ no cycle considered in the construction of h^- contained e, hence $h^-(e) = 0$ and $h(e) = h^+(e) = h_0^+(e) = g(e)$. Therefore $|h| = |g|$. Since $h \leq h^+ \leq g$, all that remains to prove is that $d^-_h(v) \leq |h|$ for every vertex v.

Since $|g|$ is finite, $\sum_{e\in E} h_i(e)$ is finite for every $i < \omega$ by the construction of the h_i. Let $x \in V$ and let X be the set of all vertices from which x is reachable via $supp(h_i)$. Note that $h_i(E(x, X-x)) = 0$, since otherwise there would exist a cycle in $supp(h_i)$. Thus, since $d^-_{h_i}(y) \geq d^+_{h_i}(y)$ for every $y \in V \setminus \{s\}$, we have

$$\sum_{y\in X-x} (d^-_{h_i}(y) - d^+_{h_i}(y)) \geq -d^+_{h_i}(s) = -|h_i|.$$

As $\sum_{e\in E} h_i(e)$ is finite, we have

$$\sum_{y\in X-x} (d^-_{h_i}(y) - d^+_{h_i}(y)) = h_i(E(V \setminus (X-x), X-x)) - h_i(E(X-x, V \setminus (X-x))).$$

By the choice of X, we have $h_i(E(V \setminus (X-x), X-x)) = 0$ and

$$h_i(E(X-x, V \setminus (X-x))) \geq h_i(E(X-x, x)) = d^-_{h_i}(x).$$

This yields $-|h_i| \leq \sum_{y\in X-x}(d^-_{h_i}(y) - d^+_{h_i}(y)) \leq -d^-_{h_i}(x)$ and hence $d^-_{h_i}(x) \leq |h_i| = |h|$. Since $h = \lim_i h_i$, we have $d^-_h(x) \leq |h|$. Since x was chosen arbitrarily, this completes the proof. □

Proof of Theorem 5.1. Easily, we may assume that $OUT(s)$ consists of one edge only. Let $\alpha = sup\{|g| : g \text{ is a flow in } \Delta\}$. If $\alpha = \infty$, we may just choose flows f_i with $|f_i| = 2^i$, and let $f = \sum \frac{f_i}{2^i}$ which, then, is a flow with $|f| = \infty$.

So assume that α is finite. Define inductively flows f_i with $|f_i| = (1-(1/2)^i)\alpha$ as follows. Let $f_0 \equiv 0$. For every $i > 0$, let g_i be a flow in $RES(\Delta, f_{i-1})$ such that $|g_i| = \frac{1}{2}(\alpha - |f_{i-1}|) = (1/2)^i$ (as $|OUT(s)| = 1$, every flow in Δ of value $|f_{i-1}| + \frac{1}{2}(\alpha - |f_{i-1}|)$ yields such a flow). By Lemma 5.4 there exists a flow $k_i \leq g_i$ such that $d^-_{k_i}(x) \leq |k_i| = |g_i|$ for every vertex x. By Lemma 5.3, $f_i := f_{i-1} \oplus k_i$ is a flow of the desired value.

By the choice of the flows k_i, the values $f_{i-1}(e)$ and $f_i(e)$ differ by at most $(1/2)^i$ for each edge e. Hence the values $f_i(e)$ converge for every e; let $f(e) = \lim_i f_i(e)$. It is easy to check that f is a flow. Further, every vertex x satisfies $d^-_f(x) \leq \alpha$. Since $|f| = \lim_i |f_i| = \alpha$, this proves the theorem. □

We shall use Theorem 5.1 twice, and in both cases we shall use it with the roles of the source and the sink reversed. Still, we chose this formulation since it is more natural.

6. Orthogonal pairs in countable networks

The main result of this section is

Theorem 6.1. *In any countable network there exists an orthogonal pair of a cut and a flow.*

By Lemma 4.14, in order to prove Theorem 6.1 it suffices to show:

Theorem 6.2. *A loose countable weighted web is linkable.*

In order to prove Theorem 1.1 for the special case of digraphs containing no infinite paths [2], or no infinite outgoing paths [4], it was possible, and useful, to reduce the problem to the special case of bipartite digraphs. Here we are going to use a similar reduction in order to deduce Theorem 6.2 from its bipartite counterpart. This is done by the following transformation.

Let $\Delta = (D, A, B, w)$ be a weighted web. We define a bipartite weighted web $\Gamma = (D', A', B', w')$ in the following way. For each vertex $v \in V(D)\setminus A$ we introduce a new vertex v_B. For each vertex $v \in V(D) \setminus B$ we introduce a new vertex v_A. We set $A' = \{v_A \mid v \in V(D) \setminus B\}$, $B' = \{v_B \mid v \in V(D) \setminus A\}$, $V(D') = A' \cup B'$, $E(D') = \{(u_A, v_B) \mid (u, v) \in E(D)\} \cup \{(v_A, v_B) \mid v \in V(D) \setminus (A \cup B)\}$, $w'(v_A) = w(v)$ for $v \in V(D) \setminus B$ and $w'(v_B) = w(v)$ for $v \in V(D) \setminus A$.

If S is a separating set in Γ and we write $A_S = \{v \mid v_A \in S\}$, $B_S = \{v \mid v_B \in S\}$ then it is straightforward to check that $S' = (A_S \cap B_S) \cup (A \cap A_S) \cup (B \cap B_S)$ is a separating set in Δ. Moreover, waves in Γ induce waves in Δ. Indeed, given a wave f in Γ with $TER(f) = S$, define the function f' on $E(D)$ by $f'(u, v) = f(u_A, v_B)$. We have:

Lemma 6.3. *f' is a wave in Δ with $TER(f') = S'$.*

Proof. Let us first prove that f' is a current. For this end, we only have to show that $d^+_{f'}(v) \le d^-_{f'}(v)$ for every $v \in V \setminus A$. This is clearly true for $v \in B$, so we may assume $v \in V \setminus (A \cup B)$. By construction, we have $d^+_{f'}(v) = d^+_f(v_A) - f(v_A, v_B)$ and $d^-_{f'}(v) = d^-_f(v_B) - f(v_A, v_B)$. Hence we are done if $d^+_f(v_A) \le d^-_f(v_B)$. So let us assume that $d^+_f(v_A) = d^-_f(v_B)$. Since $d^+_f(v_A) \le w'(v_A) = w(v) = w'(v_B)$, we have $d^-_f(v_B) < w'(v_B)$ and hence $v_B \notin TER(f)$. Likewise, we have $d^+_f(v_A) > d^-_f(v_B) \ge 0$ and hence $v_A \notin TER(f)$. Since v_A and v_B are connected by an edge, this contradicts the fact that f is a wave.

Now let us prove $TER(f') = S'$. Clearly, we have $TER(f') \cap A = A \cap A_S$ and $TER(f') \cap B = B \cap B_S$. Thus, it remains to show that $TER(f') \setminus (A \cup B) = A_S \cap B_S$. Let $v \in A_S \cap B_S$. Since $v \in A_S$, we have $v_A \in SNK(f)$ and thus $v \in SNK(f')$. Finally, we have $v_B \in SAT(f)$, which yields $v \in SAT(f')$, since $f(v_A, v_B) = 0$. Hence $TER(f') \setminus (A \cup B) \supseteq A_S \cap B_S$. Now let $v \in TER(f') \setminus (A \cup B)$. Then, $w(v) = d^-_{f'}(v) \le d^-_f(v_B) \le w'(v_B) = w(v)$, which means $v \in B_S$ and $d^-_f(v_B) = d^-_{f'}(v)$. The latter yields $f(v_A, v_B) = 0$. Since $v \in SNK(f')$, we have $v \in A_S$.

Therefore, $TER(f') = S'$. Since S' is an A-B-separator, f' is a wave in Δ. $\square$

Lemma 6.4. *If the zero wave in Γ is a hindrance, then the zero wave in Δ is also a hindrance.*

Proof. Suppose the zero wave f_0 in Γ is a hindrance and let v_A be a hindered vertex, that is, $w'(v_A) > 0 = d^+_{f_0}(v_A)$ and $v_A \notin \mathcal{E}(TER(f_0))$. In other words, every neighbour u_B of v_A lies in $TER(f_0)$ and hence satisfies $w'(u_B) = 0$. If v_B existed, we would have $w'(v_B) = w'(v_A) > 0$. Hence v_B does not exist, which means that $v \in A$. Further, every neighbour u of v in Δ satisfies $w(u) = w'(u_B) = 0$, since u_B is a neighbour of v_A in Γ. Therefore, $v \in A \setminus \mathcal{E}(TER(f'_0))$ and $w(v) = w'(v_A) = 0$. Hence the zero wave f'_0 in Δ is a hindrance. $\square$

Our next aim is to prove:

Theorem 6.5. *A countable loose bipartite weighted web is linkable.*

Theorem 6.5 implies Theorem 6.2. Indeed, Lemmas 6.3 and 6.4 imply that if Δ is loose then so is Γ. On the other hand, if f is a linkage in Γ, then the function f' defined above satisfies $d^+_{f'}(v) = w'(v_A) - f(v_A, v_B) \geq d^-_f(v_B) - f(v_A, v_B) = d^-_{f'}(v)$ for $v \in V(D) \setminus B$ and $d^+_{f'}(a) = w(a)$ for $a \in A$. Thus, applying ideas similar to those in the proof of Lemma 4.8, we can easily use f' to obtain a linkage of Δ.

The rest of this section will be devoted to the proof of Theorem 6.5. Henceforth Γ will denote a countable bipartite weighted web with sides A and B and weight function w.

Definition 6.6. If f is a current in Γ we write $\Gamma - f$ for the weighted web $(D, A, B, w - d_f)$.

Lemma 6.7. *Let u be a non-negative function on B such that $\varepsilon := \sum_{v \in B} u(v)$ is finite. Let w' be the weight function on V defined by $w' \restriction A = w \restriction A$ and $w' \restriction B = (w \restriction B) - u$. If $\Xi = (D, A, B, w')$ is $(> \varepsilon)$-hindered then Γ is hindered.*

Proof. Let f be a $(> \varepsilon)$-hindrance in Ξ, and let $a \in A \setminus \mathcal{E}(TER(f))$ be a $(> \varepsilon)$-hindered vertex for f, that is $w(a) - d_f(a) > \varepsilon$. We define a network Ψ, as follows. The vertex set of Ψ is $V(\Gamma) \cup \{t\}$, where t is a new vertex added (recovering, in fact, the sink vertex of the network from which the web Γ was obtained). The source vertex of Ψ is a, and its sink vertex is t. The edges of Ψ are all edges of Γ, taken each in both directions, together with $\{(y, t) \mid y \in B\}$. Its capacity function is defined by $c_\Psi(x, y) = \max(w(x), w(y)) + 1$, $c_\Psi(y, x) = f(x, y)$ for all $(x, y) \in E(\Gamma)$, and $c_\Psi(y, t) = u(y)$ for all $y \in B$. By Theorem 5.1 (with the roles of the source and the sink reversed) there exists in Ψ a flow j maximizing the in-degree of t, and satisfying $d^+_j(a) \leq d^-_j(t) \leq \varepsilon$. Note that for $x \in \mathcal{E}(TER(f)) \cap A$, we have $c_\Psi(e) = 0$ for each $e \in IN(x)$ and thus $d^-_j(x) = d^+_j(x) = 0$.

Call a vertex $r \in V$ *reachable* (from a) if there exists a path P from a to r in Ψ such that $c_\Psi(e) - j(e) > 0$ for all $e \in E(P)$. Note that $c_\Psi(e) - j(e) > 0$ for each A–B edge e. Hence, if a vertex in A is reachable then so are all its neighbours in B. Let g be the flow defined by letting $g(e) = 0$ if e has at least one unreachable endpoint and $g(e) = (f \oplus j)(e)$ otherwise. We shall show that g is a wave in Γ. First note that g is a current since $d_g(x) \leq d_{f \oplus j}(x) = d_f(x) + j(x, t) \leq w(x)$ for every $x \in V(\Gamma)$. Suppose, for contradiction, that $TER(g)$ is not A-B separating, in which case there exists an edge (x, y) such that neither x nor y are in $TER(g)$. Since $x \notin TER(g)$ it is reachable and so is y; indeed, if x was unreachable, we would have $x \in SNK(g)$ by definition of g, and hence $x \in TER(g)$. Thus, there exists a path P from a to y such that $c_\Psi - j$ is positive on the edges of P.

If $x \in TER(f)$, we have $d^-_j(x) = d^+_j(x) = 0$. This yields $d^+_g(x) = 0$ and thus $x \in TER(g)$, a contradiction. Thus, since f is a wave, $y \in TER(f)$. Since $y \notin TER(g)$, it is saturated by f but not by g. This means that $c_\Psi(y, t) - j(y, t) > 0$.

Thus the flow j in Ψ can be augmented along P, by adding some small number ζ on all edges of P and on (y,t). This contradicts the maximality of $d^-_j(t)$.

Therefore, g is a wave in Γ. Since $d^+_j(a) \le d^-_j(t) \le \varepsilon$, we have $d_g(a) < w_\Gamma(a)$. Thus a witnesses the fact that g is a hindrance in Γ, which proves the lemma. $\square$

Lemma 6.7 and Observation 4.6 imply:

Corollary 6.8. *If g is a current in Γ with $\sum_{v\in B} g(v) = \varepsilon$, and if $\Gamma - g$ is $(>\varepsilon)$-hindered, then Γ is hindered.*

If $\Gamma = (D, A, B, w)$ is a weighted web and g a real function on the vertices of Γ such that $g(v) \le w(v)$ for every $v \in V(D)$, we write $\Gamma - g$ for the weighted web $(D, A, B, w - g)$.

Lemma 6.9. *Let $\Omega = (D, A, B, w)$ be a loose bipartite weighted web, and let b be an element of B with $w(b) > 0$. Then there exists $\varepsilon > 0$ such that $\Omega - \varepsilon\chi(\{b\})$ is unhindered.*

Proof. Without loss of generality we may assume that $w(b) \ge 1$. This means that $\Omega - \frac{1}{n}\chi(\{b\})$ is defined for all positive integers n. Suppose, for contradiction, that $\Omega - \frac{1}{n}\chi(\{b\})$ contains a hindrance g_n for every $n = 1, 2, 3....$ Clearly, $b \in TER(g_n)$, since otherwise g_n would be a hindrance in Ω. We define a wave g_ω in $\Omega - \chi(\{b\})$ as follows. First, for every i, let $\tilde{g}_i$ be a wave in $\Omega - \chi(\{b\})$ obtained from g_i by reducing its value on some edges at b so that $d^-_{\tilde{g}_i}(b) = w(b) - 1$. Then, let $f_n = g_1 {}^\frown \tilde{g}_2 {}^\frown \ldots {}^\frown g_n$ and let $g_\omega = \sup f_n$. By Lemma 4.3, g_ω is a wave in $\Omega - \chi(\{b\})$.

Now let $h_n = g_n {}^\frown g_\omega$. It is easy to check that h_n is a wave in $\Omega - \frac{1}{n}\chi(\{b\})$, even though g_ω is not: $g_\omega \restriction ((\Omega - \frac{1}{n}\chi(\{b\}))/g_n)$ is a wave in $(\Omega - \frac{1}{n}\chi(\{b\}))/g_n$ and hence h_n is a wave, by Lemma 4.11. Let $T = TER(g_\omega) \cap B$ and let $S = A \setminus RF(T)$. Then, by Lemma 4.12, $T \supset TER(g_n) \cap B$ for all n, and hence $T = TER(h_n) \cap B$ for all n. The waves h_n all play in the same arena - the web induced on $(A \setminus S) \times T$.

Similarly with Lemma 6.7, we can define a network Ψ with sink b and source s, where s is a new vertex added, joined to all vertices in $A \setminus S$. In Ψ, we can apply Theorem 5.1 to the flows $h_2 - h_1, h_3 - h_1, \ldots$, to deduce that there exists a current k in Ψ of value 1. Then, $h_1 \oplus k$ is a current in Ω saturating all vertices in T, and is thus a non-zero wave in Ω, contradicting the fact that Ω is loose. $\square$

We shall use Lemma 6.9 for our next lemma:

Lemma 6.10. *Let $\Omega = (D, A, B, w)$ be a loose bipartite weighted web, and let a be any element of A. Then, there exists a current f such that $d_f(a) = w(a)$ and $\Omega - f$ is loose.*

Proof. We may assume that $w(a) > 0$ since otherwise we could choose $f \equiv 0$. We choose recursively vertices $y_\theta \in B$, flows f_θ and networks Ω_θ, for countable ordinals θ, as follows. Since $w(a) = 0$ and since Ω is unhindered, there exists an edge $(a, y_0) \in OUT(a)$, such that $w(y_0) > 0$. By Lemma 6.9 and Observation 4.6 we can find $\varepsilon_0 > 0$ such that $\Omega - \varepsilon_0\chi(\{a, y_0\})$ is unhindered. Let k_0 be a maximal wave in $\Omega - \varepsilon_0\chi(\{a, y_0\})$. Define $f_0 = \varepsilon_0\chi(\{(a, y_0)\}) + k_0$. Since k_0 is maximal, $\Omega - f_0$ is loose; for if $\Omega - f_0$ contained a wave g which is non-zero or a hindrance, $k_0 + g$ would be a wave in $\Omega - f_0$, contradicting either the maximality of k_0 or the fact that $\Omega - f_0$ is unhindered. Let $\Omega_1 = \Omega - f_0$. If $w_{\Omega_1}(a)$, i.e. the capacity of a in Ω_1, is greater than 0, then there exists $(a, y_1) \in OUT(a)$ with $w_{\Omega_1}(y_1) > 0$. Thus we can find $\varepsilon_1 > 0$ such that $\Omega_1 - \varepsilon_1\chi(\{a, y_1\})$ is unhindered. Taking a maximal flow k_1 in $\Omega_1 - \varepsilon_1\chi(\{a, y_1\})$ and defining $f_1 = \varepsilon_1\chi(\{(a, y_1)\}) + k_1$, the weighted web $\Omega_1 - f_1 = \Omega - f_0 - f_1$ is then loose.

We continue this way transfinitely until either the capacity of a has been reduced to 0 or we have obtained a hindered weighted web. For each ordinal α write $f_\alpha = \sum_{\theta<\alpha} f_\theta$. For successor ordinals, the currents f_α and the weighted webs Ω_α are defined as exemplified above. For limit ordinals α define $\Omega_\alpha = \Omega - f_\alpha$. We wish to show that Ω_α is unhindered for every α. By the construction, this is automatically true for successor ordinals α. Thus we only have to show:

Assertion 6.11. Ω_α *is unhindered for all limit countable ordinals* α.

The proof is by induction on α. Let α be a limit ordinal, and assume that $\Omega - f_\nu$ is unhindered for all limit ordinals $\nu < \alpha$. Clearly, hindrances cannot appear at non-limit ordinals, and thus we may assume that $\Omega - f_\nu$ is unhindered for all $\nu < \alpha$. Assume, for contradiction, that there exists a hindrance h in Ω_α. Let $z \in A$ be a hindered vertex and let $\delta = w_{\Omega_\alpha}(z) - d_h(z)$. Since $\sum_{\theta<\alpha} d_{f_\theta}(a)$ is bounded (by $w(a)$ for instance), there is some ordinal ν such that $\sum_{\nu<\theta<\alpha} d_{f_\theta}(a) < \delta$. In particular, $\sum_{\nu<\theta<\alpha} \varepsilon_\theta < \delta$. Since $f_\alpha = f_\nu + \sum_{\nu<\theta<\alpha} \varepsilon_\theta\chi(a, y_\theta) + \sum_{\nu<\theta<\alpha} k_\theta$, the current $\sum_{\nu<\theta<\alpha} k_\theta + h$ is a ($\geq \delta$)-hindrance in $\Omega - f_\nu - \sum_{\nu<\theta<\alpha} \varepsilon_\theta\chi(a, y_\theta)$. But since $\sum_{\nu<\theta<\alpha} \varepsilon_\theta < \delta$, this contradicts the fact that $\Omega - f_\nu$ is unhindered by Corollary 6.8. This proves the assertion.

Since $w_{\Omega_{\theta+1}}(a) < w_{\Omega_\theta}(a)$ for every θ, the process must stop at some countable ordinal α. But this can only happen when $w_{\Omega_\alpha}(a) = 0$. Taking $f = f_\alpha$ for α satisfying this condition yields the lemma. □

Applying this lemma recursively, we can now achieve our aim:

Proof of Theorem 6.5. Enumerate the vertices in A as $a_1, a_2, \ldots$. Applying Lemma 6.10 to Δ with $a = a_1$ we get a current f_1 in Δ saturating a_1, and having the property that $\Delta - f_1$ is loose. Using the same lemma again, we get a current f_2 in $\Delta - f_1$ saturating a_2 in this weighted web, and such that $\Delta - f_1 - f_2$ is loose. Continuing this way, we find a sequence f_i of currents, where f_i saturates a_i in $\Delta - \sum_{j<i} f_j$. The current $\sum f_i$ is then the desired linkage of Δ. □

As already mentioned, Theorem 6.5 implies Theorem 6.2, which in turn implies Theorem 6.1.

7. Mundane flows and attainability

As mentioned in the introduction, Theorem 6.1 does not generalize Theorem 1.1, since the flow is allowed to contain infinite paths. One could try to generalize Theorem 1.1 by only considering flows that do not contain infinite paths:

Definition 7.1. A flow f is *mundane* if (seen as a vector in $\mathbb{R}_+^E$) it can be written as $f = \sum_{i\in I} \theta_i \chi_E(E(P_i))$, where θ_i is a positive real number and P_i is an s-t path.

Problem 7.2. *Does there exist an orthogonal pair of a cut and a mundane flow for every infinite network?*

The results proved so far answer this question for certain networks. A *trail* in a network is a directed walk in which no edge appears more than once.

Corollary 7.3. *In every countable network $\Delta = (D, s, t, c)$ that contains no infinite trail, there is an orthogonal pair of a cut and a mundane flow.*

Proof. By the transformation of Section 3, Δ yields a weighted web $\Gamma = (D', A, B, w)$. Recall that $V(D') = E(D)$. By Theorem 6.2 and Lemma 4.14 there is an orthogonal pair of a separating set S and a web-flow f in Γ. We may assume S to be essential. We claim that there is a *mundane web-flow* $f' \leq f$ that is also orthogonal to S, where mundane web-flows are defined analogously to mundane flows.

Since Δ contains no infinite trails, there are no infinite paths in Γ; we will use this fact to construct f'. Inductively for countable ordinals i we will choose A–B paths P_i and positive real numbers θ_i so that the function $f_i := \sum_{j \leq i} \theta_j \chi(E(P_j))$ is a mundane web-flow with $f_i(e) \leq f(e)$ for each edge e. Let i be a countable ordinal and assume that P_j and θ_j have been defined for all $j < i$. Then, since each f_j satisfies $f_j \leq f$ by assumption, the function $f_{<i} := \sum_{j<i} \theta_j \chi(E(P_j))$ is a mundane web-flow with $f_{<i} \leq f$. If $f_{<i}(e) = f(e)$ for every $e \in E(A, V(D') \setminus A)$, we terminate the construction and put $f' := f_{<i}$. Otherwise, since Γ contains no infinite paths, the support of the web-flow $f - f_{<i}$ contains an A–B path P_i; let $\theta_i := \min\{f(e) - f_{<i}(e) \mid e \in E(P_i)\}$. Clearly, f_i is a mundane web-flow with $f_i \leq f$. Since $supp(f_i) \subsetneq supp(f_{<i})$ and Γ is countable, the construction terminates after countably many steps.

We thus have a mundane web-flow $f' \leq f$ that coincides with f on $E(A, V(D') \setminus A)$. We have to show that f' is orthogonal to S. Since $f' \leq f$ and f is orthogonal to S, it suffices to show that $S \subset SAT(f')$. If $d^-_{f'}(s) < w(s)$ for a vertex $s \in S \setminus A$, then $d^-_{f-f'}(s) > 0$. Since $f - f'$ is a web-flow, no vertex in the digraph $\tilde{D} = (V(D'), supp(f - f'))$ that does not lie in $A \cup B$ has degree 1. Hence s lies on an A–B path in $\tilde{D}$, or on an infinite path, or on a cycle. By the choice of f', there are no A–B paths in $\tilde{D}$, and $\tilde{D}$ does not contain infinite paths since D' does not. So s lies on a cycle C in $\tilde{D}$, which is clearly also a cycle in $supp(f)$. Since S is essential we have $s \in RF(S) \setminus RF^\circ(S)$, and hence C contains an edge e from $V(D') \setminus RF^\circ(S)$ to $RF(S)$. But then $e \in supp(f)$ and thus $f(e) > 0$, contradicting the fact that f is orthogonal to S.

We have shown that there is an orthogonal pair of a separating set S and a mundane web-flow $f' = \sum_{i \in I} \theta_i \chi(E(P_i))$ in Γ. These pair can easily be translated into an orthogonal pair of a cut F and a mundane flow g in Δ: The vertex set S in D' is an edge set in D and it is s–t separating in D since it is A–B separating in D'; hence it contains a cut F in D. Every A–B path P_i in D' corresponds to an s–t trail P'_i in D; let g' be the function on $E(D)$ defined by $g' := \sum_{i \in I} \theta_i \chi(E(P'_i))$. It is easy to see that g' is a flow in Δ orthogonal to S and hence also to F. Therefore, each P'_i meets F in precisely one edge. Every P'_i contains an s–t path Q_i; let $g := \sum_{i \in I} \theta_i \chi(E(Q_i))$. Then Q_i meets F at the same edge as P'_i does, and hence g is a mundane flow in Δ orthogonal to the cut F. $\square$

In the remainder of this section we show that the infimum σ of the capacities of the s–t cuts in a network equals the supremum τ_m of the values of the mundane flows. Moreover, we show that σ is attained by some cut but τ_m need not be attained by any mundane flow.

Definition 7.4. Given a countable network $\Delta = (D, c, s, t)$, let

$$\sigma := \inf\{c[F] : F \text{ is an } s\text{–}t \text{ cut}\}, \text{ and}$$
$$\tau_m := \sup\{|f| : f \text{ is a mundane flow in } \Delta\}.$$

Theorem 7.5. *Let $\Delta = (D, c, s, t)$ be a countable network. The following statements hold:*

(i) *Δ has an s–t cut F of minimal capacity σ, and*
(ii) *$\tau_m = \sigma$.*

Proof. For every positive integer i, let c_i be the function obtained from c by cutting off everything behind the ith decimal; formally, $c_i(e) = \lfloor 10^i c(e) \rfloor / 10^i$. In the network $\Delta_i = (D, c_i, s, t)$, all capacities are multiples of 10^{-i}, hence we can use Theorem 1.1 to find an orthogonal pair of a mundane flow f_i and a cut F_i in Δ_i. Since $c_i \leq c$, f_i is also a flow in Δ. This yields $c_i[F_i] = |f_i| \leq \tau_m$. We will use the cuts F_i to construct a cut F with capacity τ_m.

First, enumerate all edges in $E(D)$ as $e_1, e_2, \ldots$. Then, inductively for every positive integer i, if there is an integer m such that $m > j_l$ for all $l < i$ and the set $\{e_{j_1}, \ldots, e_{j_{i-1}}, m\}$ is contained in infinitely many of the cuts $F_1, F_2, \ldots$, then let j_i be the smallest such integer m. If no such m exists, then stop.

If j_i exists for all i, we end up with a set $F' = \{e_{j_1}, e_{j_2}, \ldots\}$ of edges. Now choose a subsequence of $F_1, F_2, \ldots$ as follows: For every positive integer i, let k_i be the smallest integer such that $k_i > k_l$ for all $l < i$ and the set $\{e_{j_1}, \ldots, e_{j_i}\}$ is contained in F_{k_i}.

If for some i there is no j_i as desired, we end up with a finite set $F' = \{e_{j_1}, \ldots, e_{j_{i-1}}\}$ and we choose $F_{k_1}, F_{k_2}, \ldots$ to be the subsequence of $F_1, F_2, \ldots$ consisting of all cuts that contain F'.

In both cases, every edge e_l that is contained in infinitely many of the cuts $F_{k_1}, F_{k_2}, \ldots$ is contained in F', since it must have been chosen as e_{j_i} at some step i. We claim that $c[F'] \leq \tau_m$. Indeed, for every $\varepsilon > 0$, there is a finite subset F'' of F' with $c[F''] \geq c[F'] - \frac{1}{2}\varepsilon$. For sufficiently large i, we have $c_i[F''] \geq c[F''] - \frac{1}{2}\varepsilon$ and thus $c[F'] \leq c_i[F''] + \varepsilon \leq \tau_m + \varepsilon$. With $\varepsilon \to 0$, this yields $c[F'] \leq \tau_m$. We further claim that F' separates s from t. Indeed, let P be an s–t path. Since P is finite and $F_{k_1}, F_{k_2}, \ldots$ infinite, P contains an edge that is contained in infinitely many of the cuts $F_{k_1}, F_{k_2}, \ldots$, and is thus contained in F', so F' meets every s–t path. Therefore, F' contains a cut F which, then, satisfies $c[F] \leq \tau_m$.

This shows that $\sigma \leq c[F] \leq \tau_m$. Combining with the trivial inequality $\tau_m \leq \sigma$ we obtain the required result. $\square$

The remaining question is whether there is always a mundane flow of value τ_m. The following example shows that this is not the case, providing a negative answer to Problem 7.2.

Example 7.6. We construct a locally finite network in which there is no mundane flow of maximal value. We start with a disjoint union of (directed) paths $Q_i = x_0^i x_1^i x_2^i x_3^i$, $i = 1, 2, \ldots$. For every positive integer k, let each edge e on any path Q_i with $2^{k-1} \leq i \leq 2^k - 1$ have capacity $c(e) = 1/2^k$. Further, for each such k and i, we attach the paths Q_{2i} and Q_{2i+1} to Q_i by adding the edges (x_0^i, x_0^{2i}), (x_3^{2i}, x_2^i) (to attach Q_{2i}), (x_1^i, x_0^{2i+1}), and (x_3^{2i+1}, x_3^i) (to attach Q_{2i+1}). Let each such edge e have capacity $c(e) = 1/2^k$. We denote the resulting digraph by D. The definition of the network $\Delta = (D, c, s, t)$ is completed by choosing $s = x_0^1$ and $t = x_3^1$ (see Figure 7.1).

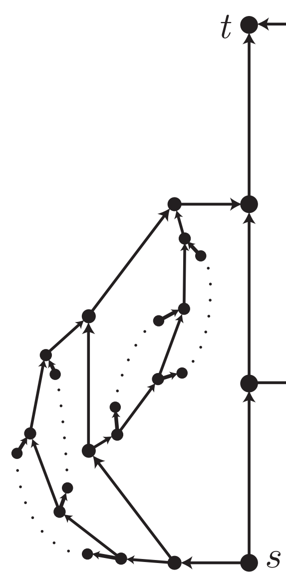

FIGURE 7.1. A locally finite network with no mundane flow of maximal value

Clearly, D is locally finite (in fact it has maximum degree 3). For every positive integer k, there exists a mundane flow of value $1 - 1/2^k$: It is easy to see that for each positive integer i, there is exactly one s–t path that contains Q_i; denote it by P_i. Then $f_k := \sum_{i=1}^{2^k-1} \frac{1}{2^k} P_i$ is a mundane flow of value $1 - 1/2^k$. This shows that $\tau_m \geq 1$.

We claim that there is no mundane flow in Δ that has value 1. Indeed, suppose for contradiction that f is a mundane flow with $|f| = 1$. Let $e := (x_1^1, x_2^1)$ and $d := (x_1^1, x_0^3)$. Applying Kirchhoff's first law to x_1^1 we obtain $f(d) \leq 1/2 - f(e)$. However, since $F = \{d, (x_2^1, x_3^1)\}$ is an s–t cut with $c[F] = 1$, f must saturate F and thus $f(d) = 1/2$ whence $f(e) = 0$ holds. Similarly, we can prove that $f(g) = 0$ holds for every edge g of the form (x_1^i, x_2^i). Since these edges form an s–t cut we obtain a contradiction to the fact that f is mundane.

8. FLOWING THROUGH AN END

In this section we consider constraints on flows that are weaker than being mundane, in order to allow for flows to flow, in a sense, through ends of the digraph. As an example look at the flows in Figure 8.2 and Figure 8.3. The definition of a mundane flow does not distinguish between the two and rejects both. However, there is an important difference: The flow in Figure 8.2 disappears in the left end of the graph and comes back from the right one, while the flow in Figure 8.3 just flows through the left end. In this section we study flows of the second kind. In order to distinguish them formally from other flows we need an analog of Kirchhoff's first law for ends. In the case of Figure 8.3 it is possible to say how much flow arrives at the left end and how much flow leaves it, but in general this is not possible: look for example at the network in Figure 7.1. The flows f_k used there have a limit flow g. Now for every ray R in this network the values of g along R converge to 0, however there is some flow running to infinity and coming back. Similarly to the examples in Figure 8.2 and Figure 8.3, it is possible to construct flows like g where the flow does flow out of the same ends it flows in (like in Figure 8.3 and Figure 7.1) or it does not (like in Figure 8.2). For flows like g it is not clear how to make precise the assertion than the ends satisfy Kirchhoff's first law. The following definition accomplishes this task in an elegant way:

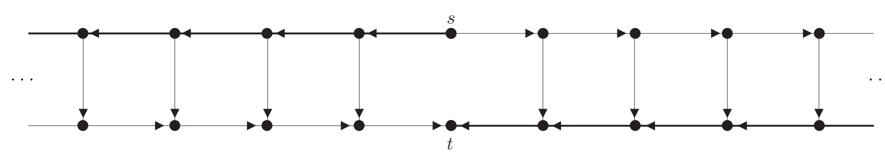

FIGURE 8.2. A network and a flow. Thick edges carry a flow of value 1; thin edges carry no flow. This flow flows into the left end of the graph and returns through the right one.

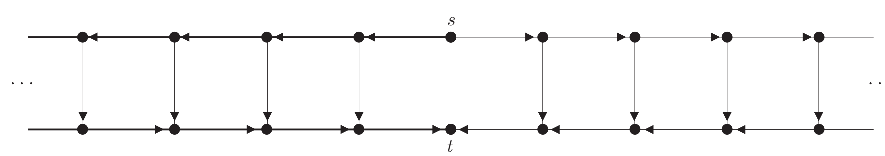

FIGURE 8.3. A flow flowing through the left end of the graph.

Definition 8.1. We will call a flow in a network $\Delta = (D, c, s, t)$ *finite-cut-respecting* if for every cut $E(S, T)$ in D (where $T = V(D) \setminus S$) with $s \in S$ that consists of finitely many edges we have

$$f[E(S,T)] = \begin{cases} f[E(T,S)] & \text{if } t \in S, \\ f[E(T,S)] + |f| & \text{if } t \in T. \end{cases} \tag{1}$$

Let $\tau_w := sup\{|f| : f \text{ is a finite-cut-respecting flow}\}$, and let σ_w be the infimum of the capacities of all s–t cuts consisting of finitely many edges.

To see why this definition can be thought of as an analog of Kirchhoff's first law for ends note that in a locally finite network a cut consisting of finitely many edges cannot separate two rays in the same end. It is easy to check that g as well as the flow in Figure 8.3 is finite-cut-respecting while the flow in Figure 8.2 is not.

Theorem 8.2. *In every locally finite network $\Delta = (D, c, s, t)$, $\sigma_w = \tau_w$ holds. Moreover, there is a finite-cut-respecting flow f such that $|f| = \tau_w$.*

Proof. For every edge e in D let I_e be the real interval $[0, c(e)]$, and define the topological space $X := \Pi_{e \in E(D)} I_e$. By Tychonoff's theorem X is compact.

Pick an s–t path P in D, and for every $i \in \mathbb{N}$, let $\Delta_i = (D_i, c_i, s, t)$ be the finite network obtained from Δ by contracting each component C of $D - \{x \in V(D) \mid d(x, P) \leq n\}$ to a vertex v_C, and letting $c_i(e) = c(e)$ for every edge in this network.

By the MFMC theorem (for finite networks) there is a flow f_i in Δ_i such that $|f_i| = \sigma_i$, where σ_i denotes the minimum capacity of an s–t cut in Δ_i. For every n, f_i corresponds to a point x_i in X: the point that has value $f_i(e)$ at the coordinate I_e of X for every $e \in E(D_i)$ and value 0 at every other coordinate. Since X is compact, the sequence $x_1, x_2, \dots$ has an accumulation point x, which determines a function $f : E(D) \to \mathbb{R}$.

We claim that f is a finite-cut-respecting flow in Δ; indeed, if (1) is violated by f for some finite cut B, in particular if Kirchhoff's law is violated at some vertex, then there is a basic open neighbourhood $O \ni x$ in X, chosen by taking a small enough interval of I_e around $f(e)$ for every $e \in B$, such that every function in O also violates (1) at B. But this cannot be the case since any such O contains some x_i where i is large enough so that B is a cut in D_i.

Similarly, it is not hard to check that $|f|$ is an accumulation point of the sequence $\{|f_i|\}_{i\in\mathbb{N}}$. Since any cut in some D_i is also a cut in D, we have $\sigma_i \geq \sigma_w$, and since $|f_i| = \sigma_i$, we obtain $|f| \geq \sigma_w$. But $|f| \leq \tau_w \leq \sigma_w$ by (1), thus $|f| = \tau_w = \sigma_w$ □

Thus the value τ_w is always attained by some finite-cut-respecting flow. However, σ_w does not have to be attained by some finite cut, as shown by the following example.

Example 8.3. Starting with the network of Example 7.6, we modify the capacities of its edges as follows. For every edge e that is the middle edge (x_1^i, x_2^i) of some path Q_i let $c'(e) = 0$; for every other edge f, if $c(f) = 1/2^k$ then let $c'(f) = 1/4^k$. Now the resulting network $\Delta' = (D, c', s, t)$ has $\sigma_w = 0$ but the only cut of capacity 0 is the infinite cut consisting of all the middle edges of the Q_i.

Although the definition of a finite-cut-respecting flow allows flows through ends and forbids flows like the one in Figure 8.2, there are also instances of finite-cut-respecting flows that may seem unnatural. Look for example at Figure 8.4; it shows a finite-cut-respecting flow of value 1 from s to t, in a network that contains no finite directed s–t path. The following definition bans such flows.

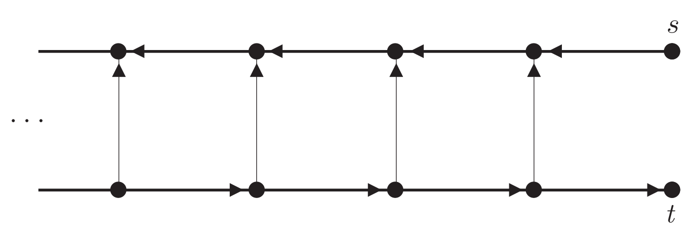

Figure 8.4. A non-zero finite-cut-respecting flow in a network with no finite directed s–t path.

Definition 8.4. We will call a flow f in a network $\Delta = (D, c, s, t)$ *cut-respecting* if it is finite-cut-respecting and moreover for every s–t cut $E(S,T)$ in D we have

$$
\begin{aligned}
|f| + f[E(T,S)] &\leq c[E(S,T)], \text{ and} \\
f[E(S,T)] &\leq c[E(T,S)] + |f|.
\end{aligned}
\tag{2}
$$

Intuitively, the first condition demands that if some flow circumvents an infinite s–t cut $E(S,T)$, then this circumvention does not exceed the amount that could flow through $E(S,T)$ given its capacity $c[E(S,T)]$, taking into account that the flow through $E(S,T)$ should also compensate for any flow $f[E(T,S)]$ in the inverse direction. The second condition demands that if some s–t cut carries more flow than $|f|$, then the excess is not greater than the amount than could go back through the inverse cut.

Let $\tau_s := sup\{|f| : f \text{ is a cut-respecting flow}\}$.

Theorem 8.5. *In every locally finite network $\Delta = (D, c, s, t)$ we have $\sigma = \tau_s$. Moreover, there is a cut-respecting flow f such that $|f| = \tau_s$ and an s–t cut F with $c[F] = \sigma$ orthogonal to f.*

Proof. Since, clearly, every mundane flow is cut-respecting, we have $\tau_s \geq \tau_m$, and thus, by Theorem 7.5 and condition (2), $\sigma = \tau_s$. Let $f_1, f_2, \ldots$ be a sequence of mundane flows in Δ whose values converge to $\tau_m = \tau_s$. As in the proof of Theorem 8.2, for every edge e in D let I_e be the real interval $[0, c(e)]$, and define the topological space $X := \Pi_{e \in E(D)} I_e$. Every f_i corresponds to a point x_i in X: the point that has value $f_i(e)$ at the coordinate I_e of X for every $e \in E(D)$. Since X is compact, the sequence $x_1, x_2, \ldots$ has an accumulation point x, which determines a function $f : E(D) \to \mathbb{R}$. Similarly with the proof of Theorem 8.2, it is straightforward to check that f is a cut-respecting flow since every f_i is, and that $|f| = \tau_s$.

Let F be an s–t cut with $c[F] = \sigma$, which exists by Theorem 7.5. We claim that f saturates F. Suppose for contradiction that there is an edge $e \in F$ such that $f(e) < c(e) - \epsilon$ for some $\epsilon > 0$. Then, there is an infinite subsequence (f'_i) of (f_i) with $f'_i(e) < c(e) - \epsilon$. But this means that the f'_i are mundane flows in the network Δ' obtained from Δ by reducing $c(e)$ by ϵ. Thus, $\lim |f'_i| \leq \tau_m - \epsilon$ by Theorem 7.5 since F is a cut of capacity $\sigma - \epsilon$ in that network. This contradicts the choice of (f_i), so f saturates F as claimed. Similarly, it is easy to show that for every T–S edge e we have $f(e) = 0$, which proves that f and F form an orthogonal pair. □

It is possible to consider networks where the source s or sink t or both are ends of the underlying digraph D instead of vertices. An s–t flow of value m is, then, a function f on $E(D)$ such that $KIR(f) = V(D)$ and moreover, for every finite cut $E(S,T)$ such that s *lives* in S we have $f(E(S,T)) = m$ unless t also lives in S, in which case we have $f(E(S,T)) = 0$. Here, we say that an end lives in S if one of its rays, and thus, since $E(S,T)$ is finite, a subray of any of its rays, is contained in S; we also say that the vertices of S live in S. The interested reader will be able to confirm that the results of this section carry over to such networks and flows.

References

[1] R. Diestel. *Graph Theory* (3rd edition). Springer-Verlag, 2005.
Electronic edition available at:
`http://www.math.uni-hamburg.de/home/diestel/books/graph.theory`.

[2] R. Aharoni, Menger's theorem for graphs containing no infinite paths, *Eur. J. Combin.*, (1983)

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

[4] 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.

[5] R. Aharoni and E. Berger, Menger's theorem for infinite graphs, *submitted for publication*.

[6] K.P. Podewski and K. Steffens, Injective choice functions for countable families, *J. Combin. Th. Ser. B* **21** (1976), 40–46.

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

Department of Mathematics, Haifa University, Israel 32000
E-mail address, Eli Berger: `berger@cri.haifa.ac.il`

Universität Hamburg
E-mail address, Angelos Georgakopoulos: `georgakopoulos@math.uni-hamburg`

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

Universität Hamburg
E-mail address, Philipp Sprüssel: `spruessel@math.uni-hamburg`