An algorithm for packing connectors

J. Keijsper
University of Amsterdam
Department of Mathematics
Plantage Muidergracht 24
1018 TV Amsterdam

Abstract

Given an undirected graph $G = (V, E)$ and a partition $\{S, T\}$ of V, an S–T *connector* is a set of edges $F \subseteq E$ such that every component of the subgraph (V, F) intersects both S and T.

If either S or T is a singleton, then an S–T connector is a spanning subgraph of G. On the other hand, if G is bipartite with colour classes S and T, then an S–T connector is an edge cover of G (a set of edges covering all vertices).

An S–T connector is a common spanning set of two graphic matroids on E.

We prove a theorem on packing common spanning sets of certain matroids, generalizing a result of Davies and McDiarmid on strongly base orderable matroids.

As a corollary, we obtain an $O(\tau(n,m) + nm)$ time algorithm for finding a maximum number of S–T connectors, where $\tau(n,m)$ denotes the complexity of finding a maximum number of edge disjoint spanning trees in a graph on n vertices and m edges. Since the best known bound for $\tau(n,m)$ is $O(nm \log(m/n))$, this bound for packing S–T connectors is as good as the current bound for packing spanning trees.

1 Introduction

Let $G = (V, E)$ be an undirected graph, S a subset of its vertices, and T the complement of S in V. An S–T *connector* in G is a set F of edges such that every component of the subgraph (V, F) intersects both S and T.

Let G_S be the graph obtained from G by shrinking the set S into a vertex s, and let G_T be the graph obtained from G by shrinking T into a vertex t. The definition implies that $F \subseteq E$ is an S–T connector in G if and only if F

contains a spanning tree of G_S and a spanning tree of G_T. Here, the edges of G are identified with the edges (possibly loops) of G_S and with the edges of G_T in the obvious way. It follows that an S–T connector is a common spanning set of two matroids on E, namely the cycle matroids of the graphs G_S and G_T. Recall that a spanning set in a matroid is a set containing a base. (For matroid theory we refer to [8].)

Let k denote a nonnegative integer. It follows from the above that if G contains k edge-disjoint S–T connectors, then both G_S and G_T contain k edge-disjoint spanning trees. In this paper, we prove that the converse also holds.

Theorem 1 *G contains k edge-disjoint S–T connectors if and only if both G_S and G_T contain k edge-disjoint spanning trees.*

The proof of this theorem provides an $O(nm)$ reduction of the problem of finding a maximum number of edge-disjoint S–T connectors in a graph to the problem of finding a maximum number of edge-disjoint spanning trees in two smaller graphs. Here, n denotes the number of vertices and m the number of edges of the graph.

Tutte [7] and Nash-Williams [6] characterized the graphs with k edge-disjoint spanning trees.

Lemma 1 (Tutte, Nash-Williams) *Let $G = (V, E)$ be an undirected graph. Then G contains k edge-disjoint spanning trees if and only if $|\delta(P)| \geq k(|P| - 1)$ for every partition P of V into nonempty subsets.* $\square$

Here, if $P = \{U_1, \ldots, U_t\}$ is a (sub)partition of V (a subpartition is a collection of pairwise disjoint nonempty subsets of V), then $\delta(P) := \bigcup_{i=1}^{t} \delta(U_i)$, where $\delta(U_i)$ denotes the set of edges with one end in U_i and one end in $V \setminus U_i$.

By Lemma 1, Theorem 1 is equivalent to the following characterization.

Theorem 2 *G contains k edge-disjoint S–T connectors if and only if $|\delta(W)| \geq k|W|$ for every subpartition W of S or T.*

For a short proof and a polyhedral interpretation of Theorem 2 the reader is referred to [4]. The proof of Theorem 1 presented in this paper is more elementary and it yields an efficient algorithm for packing S–T connectors (using a subroutine for packing spanning trees as a black box). Moreover, Theorem 1 is derived as a special case of a packing result for matroids, which is interesting in its own right.

Theorem 2 has two important special cases. First, if G is bipartite with colour classes S and T, then an S–T connector is nothing but an edge cover of G (a set of edges covering all vertices), and Theorem 2 specializes to a theorem of König [5] and Gupta [3], saying that the maximum number of edge-disjoint edge covers in a bipartite graph is equal to the minimum vertex degree. Second, if either S or T is a singleton, then an S–T connector is a set of edges containing a spanning tree of G, and Theorem 2 specializes to Lemma 1.

Note that Lemma 1 is a special case of the well-known matroid base packing theorem. However, Theorems 1 and 2 do not follow from any known matroid base packing results.

2 Packing common spanning sets

Since S–T connectors are common spanning sets, matroid intersection provides a min-max relation for the minimum cardinality (or weight) of an S–T connector in G. However, for the packing of common spanning sets of two matroids no general theorem is known. In the case of strongly base orderable matroids, there is a theorem by Davies and McDiarmid [1].

Definition 1 A matroid $\mathcal{M}$ is said to be *strongly base orderable* if for every two bases B_1 and B_2 in $\mathcal{B}$ there exists a bijection $\pi : B_1 \to B_2$ such that for every subset A of B_1 the set $(B_1 \setminus A) \cup \pi(A)$ is a base of $\mathcal{M}$. (It follows that for such a π, also $(B_2 \setminus \pi(A)) \cup A$ is a base of $\mathcal{M}$.)

Theorem 3 (Davies and McDiarmid) *Let $\mathcal{M}$ and $\mathcal{N}$ be strongly base orderable matroids on the same set E. Suppose that both $\mathcal{M}$ and $\mathcal{N}$ have k pairwise disjoint bases. Then E contains k pairwise disjoint common spanning sets of $\mathcal{M}$ and $\mathcal{N}$.*

□

Graphic matroids are generally not strongly base orderable: the reader may verify that the cycle matroid of K_4 is not strongly base orderable. In fact, if $\mathcal{M}$ is the cycle matroid of K_4 and $\mathcal{N}$ is the matroid with independent sets all sets of edges of K_4 which are pairwise intersecting, then the edge set of K_4 can be partitioned into two bases of $\mathcal{M}$ and into two bases of $\mathcal{N}$ but does not contain two disjoint common spanning sets of $\mathcal{M}$ and $\mathcal{N}$.

To be able to formulate a generalization of Theorem 3 that can be applied to graphic matroids, we need the following definition.

Definition 2 Let $\mathcal{M}$ be matroid on E and let $E_0 \subseteq E$. Then E_0 is said to be an *sbo-set* for $\mathcal{M}$ if for every two bases B_1 and B_2 of $\mathcal{M}$, there exists an injective function $\pi : B_1 \cap E_0 \to B_2$, such that for every subset A of $B_1 \cap E_0$ with $|\pi(A) \setminus E_0| \leq 1$, the sets $(B_1 \setminus A) \cup \pi(A)$ and $(B_2 \setminus \pi(A)) \cup A$ are bases of $\mathcal{M}$.

It is not difficult to see that such a function π must be the identity on $B_1 \cap B_2 \cap E_0$.

The property defined in Definition 2 is preserved under taking duals and minors in the following sense: if E_0 is an sbo-set for $\mathcal{M}$, then E_0 is also an sbo-set for the dual $\mathcal{M}^*$ of $\mathcal{M}$, and $E_0 \cap E'$ is an sbo-set for the restriction of $\mathcal{M}$ to E', for any subset E' of E. It follows that $E_0 \cap E'$ is an sbo-set for any minor $\mathcal{M}'$ of $\mathcal{M}$ with ground set $E' \subseteq E$.

Also, if E_0 is an sbo-set for $\mathcal{M}$, then any subset of E_0 is an sbo-set for $\mathcal{M}$. Moreover, every singleton subset of E is an sbo-set for $\mathcal{M}$.

Note that $\mathcal{M}$ is strongly base orderable if and only if $E_0 = E$ is an sbo-set for $\mathcal{M}$. The following theorem therefore implies Theorem 3 (take $E_1 = E_2 = E = E_0$).

Theorem 4 *Let $\mathcal{M}$ be a matroid on E_1 and let $\mathcal{N}$ be a matroid on E_2. Let $E := E_1 \cup E_2$ and $E_0 := E_1 \cap E_2$. View $\mathcal{M}$ and $\mathcal{N}$ as matroids on E (consider elements of $E_2 \setminus E_1$ as loops of $\mathcal{M}$ and elements of $E_1 \setminus E_2$ as loops of $\mathcal{N}$). Suppose that E_0 is an sbo-set for $\mathcal{M}$ as well as for $\mathcal{N}$, and suppose that both $\mathcal{M}$ and $\mathcal{N}$ have k pairwise disjoint bases. Then E contains k pairwise disjoint common spanning sets of $\mathcal{M}$ and $\mathcal{N}$.*

Proof. Let $M_1, \ldots, M_k$ be k disjoint bases of $\mathcal{M}$ and let $N_1, \ldots, N_k$ be k disjoint bases of $\mathcal{N}$, such that

$$\mu := \sum_{i \neq j} |N_i \cap M_j|$$

is minimal.

Suppose $\mu > 0$. Then there are indices r and b, $r \neq b$, with $M_r \cap N_b \neq \emptyset$, say $e_0 \in M_r \cap N_b$. Observe that $e_0 \in E_0$.

We assume henceforth that $E_0 = E_0 \cap (M_r \cup M_b \cup N_r \cup N_b)$. Since any subset of an sbo-set is an sbo-set, this causes no loss of generality.

As E_0 is an sbo-set for $\mathcal{M}$, there is an injection

$$\pi : M_r \cap E_0 \to M_b$$

with the property that for every $A \subseteq M_r \cap E_0$ with $|\pi(A) \setminus E_0| \leq 1$, both $(M_r \setminus A) \cup \pi(A)$ and $(M_b \setminus \pi(A)) \cup A$ are bases of $\mathcal{M}$. For $\mathcal{N}$, there is a similar injection

$$\pi' : N_b \cap E_0 \to N_r.$$

Let D be the directed graph with vertex set E and arc set the set of pairs (e, f) from $E \times E$ with $e \in E_0$ and $f = \pi(e)$ or $f = \pi'(e)$. Moreover, define

$$Y := (N_r \setminus M_r) \cup (M_b \setminus N_b).$$

Then there exists a directed e_0–Y path in D. Indeed, $d^+_D(e) \geq 1$ for each $e \in E_0 \setminus Y$ by definition of π and π', $d^-_D(e) \leq 1$ for each $e \in E_0 \setminus Y$ because π and π' are injective and their ranges in $E_0 \setminus Y$ are disjoint, and $d^-_D(e_0) = 0$, since $e_0 \in M_r \cap N_b$. Here, $d^+_D(e)$ denotes the outdegree and $d^-_D(e)$ denotes the indegree of a vertex e of D.

Let $P = (e_0, e_1, \ldots, e_n)$ be a shortest e_0–Y path. Then $n \geq 1$. Define

$$A := \{e_i \mid 0 \leq i < n,\ i \text{ odd}\}$$

$$B := \{e_i \mid 0 \leq i < n,\ i \text{ even}\}$$

By definition either $e_1 = \pi(e_0)$ or $e_1 = \pi'(e_0)$. Without loss of generality assume $e_1 = \pi'(e_0) \in N_r$. Then $A \subseteq M_r \cap E_0$ and $B \subseteq N_b \cap E_0$. So by definition of π, $M'_r := (M_r \setminus A) \cup \pi(A)$ and $M'_b := (M_b \setminus \pi(A)) \cup A$ are bases of $\mathcal{M}$. Similarly, $N'_b := (N_b \setminus B) \cup \pi'(B)$ and $N'_r := (N_r \setminus \pi'(B)) \cup B$ are bases of $\mathcal{N}$ (note that $\pi(A)$ and $\pi'(B)$ are contained in $E_0 \cup \{e_n\}$, so $|\pi(A) \setminus E_0| \leq 1$ and $|\pi'(B) \setminus E_0| \leq 1$).

Now if we set $M'_j = M_j$ and $N'_j = N_j$ for $j \neq r, b$, then

$$(1) \qquad \bigcup_{i \neq j} (M'_i \cap N'_j) \subseteq \bigcup_{i \neq j} (M_i \cap N_j) \setminus \{e_0\}.$$

Indeed, by construction, $e_0 \in (M_r \cap N_b) \cap (M'_r \cap N'_r)$, $e_i \in (M_r \cap N_r) \cap (M'_b \cap N'_b)$ for every i with $0 < i < n$ and i odd, and $e_i \in (M_b \cap N_b) \cap (M'_r \cap N'_r)$ for every i with $0 < i < n$ and i even. Moreover, either n is odd and $e_n \in (N_r \setminus M_r) \cap (N'_b \setminus M'_r)$ or n is even and $e_n \in (M_b \setminus N_b) \cap (M'_r \setminus N'_b)$ (so $e_n \notin E_0$, or $e_n \in (M_j \cap N_r) \cap (M'_j \cap N'_b)$ for some $j \neq r$, or $e_n \in (M_b \cap N_j) \cap (M'_r \cap N'_j)$ for some $j \neq b$). Edges not on P are in M'_j whenever they are in M_j and in N'_j whenever they are in N_j, for any j.

From (1) it follows that

$$\mu' = \sum_{i \neq j} |M_i' \cap N_j'| < \sum_{i \neq j} |M_i \cap N_j| = \mu,$$

contradicting the minimality of μ.

Therefore $\mu = 0$ and the sets $S_i := M_i \cup N_i$ are disjoint common spanning sets of $\mathcal{M}$ and $\mathcal{N}$. $\square$

Any matroid partition algorithm (see for example [2]) can be used to obtain the disjoint bases $M_1, \ldots, M_k$ of $\mathcal{M}$ and the disjoint bases $N_1, \ldots, N_k$ of $\mathcal{N}$ if they exist. The above proof suggests an algorithm for modifying the M_i and the N_i in steps, such that after every step the M_i are disjoint bases of $\mathcal{M}$ and the N_i are disjoint bases of $\mathcal{N}$, and moreover the common spanning sets $M_i \cup N_i$ are disjoint after at most μ steps. Each step consists of finding a path in D and 'recolouring' the elements of E that correspond to vertices of the path. The vertices of the path can be found efficiently provided that there are subroutines available that compute $\pi(e)$ for all $e \in A$ and $\pi'(e)$ for all $e \in B$ in polynomial time, for any $A \subseteq M_r \cap E_0$ with $|\pi(A) \setminus E_0| \leq 1$ and $B \subseteq N_b \cap E_0$ with $|\pi(B) \setminus E_0| \leq 1$. (Note that the sets A and B are not input, but starting from $e_0 \in B$ and computing π'- and π-values alternatingly, an element of A or B is identified at the moment we want to compute its π- or π'-value.)

As a corollary of Theorem 4 we can derive Theorem 1. For this we need the following lemma. We write $\delta(v)$ instead of $\delta(\{v\})$.

Lemma 2 *Let $G = (V, E)$ be a graph and let $s \in V$. Then $\delta(s)$ is an sbo-set for the cycle matroid $\mathcal{M}(G)$ of G.*

Proof. Let T_r and T_b be the edge sets of two spanning trees in G. By contracting the edges in $T_r \cap T_b$, we may assume that they are disjoint.

For any edge e of $T_r \cap \delta(s)$ the *e-branch* of T_r is the component of $(V, T_r) - \{s\}$ incident with e. Define a function $\pi : T_r \cap \delta(s) \to T_b$ as follows: if $e = \{s, v\} \in T_r \cap \delta(s)$, let $\pi(e)$ be the first edge of the unique path from v to s in T_b that leaves the e-branch of T_r. Then clearly, π is injective, since different edges in $T_r \cap \delta(s)$ define different branches of T_r, and an edge of T_b is traversed in the same direction by every v–s path in T_b that uses it.

Let $A \subseteq T_r \cap \delta(s)$ with $|\pi(A) \setminus \delta(s)| \leq 1$. It has to be shown that

$$T_r' := (T_r \setminus A) \cup \pi(A), \text{ and } T_b' := (T_b \setminus \pi(A)) \cup A$$

are spanning trees. Because π is injective, T_r' and T_b' have the right cardinality.

Suppose T_r' is not the edge set of a spanning tree. Then there exists a set of vertices $U \subseteq V$ with $s \notin U$ such that $T_r' \cap \delta(U) = \emptyset$. Consequently, $T_r \cap \delta(U) \subseteq A \subseteq \delta(s)$ (so every branch of T_r is contained in U or in its complement) and $\pi(A) \cap \delta(U) = \emptyset$. Hence, for every $e \in A \cap \delta(U)$, $\pi(e)$ is contained in U. Since π is an injection and $|\pi(A) \setminus \delta(s)| \leq 1$, there is at most one edge in $A \cap \delta(U)$. It is impossible to have exactly one edge $e \in A \cap \delta(U)$, because $\pi(e) \subseteq U$ should connect two different branches of T_r, each contained in U and defined by an edge of $A \cap \delta(U)$. So $A \cap \delta(U) = \emptyset$ and hence $T_r \cap \delta(U) = \emptyset$, contradicting the fact that T_r is a spanning tree.

Finally, suppose that T_b' is not the edge set of a spanning tree. Then there is a cut $\delta(U)$ ($s \notin U$) with $T_b' \cap \delta(U) = \emptyset$. It follows that $T_b \cap \delta(U) \subseteq \pi(A)$ and that $A \cap \delta(U) = \emptyset$. Since T_b is a spanning tree, $T_b \cap \delta(U) \neq \emptyset$. Let $e \in T_b \cap \delta(U)$. Then $e = \pi(f)$ for some $f \in A \subseteq \delta(s) \setminus \delta(U)$. Say $f = \{s, v\}$. By definition of π, the path P from v to s in T_b traverses e. P intersects $\delta(U)$ in an even number of edges. So, there must be another edge $e' \neq e$ in $T_b \cap \delta(U) \subseteq \pi(A)$ on P. Since $|\pi(A) \setminus \delta(s)| \leq 1$, one of the edges e and e', say e, is in $\delta(s)$. Then e' is before e (going from v to s) on P and $e' \in \pi(A) \setminus \delta(s)$. Therefore, e' connects two different branches of T_r. On the other hand, by definition of π, every edge on P before $e = \pi(f)$ has both ends in the same branch of T_r. From this contradiction we derive that T_b' is the edge set of a spanning tree, as required. $\square$

For the function π defined in the proof of the above lemma, it is possible to compute $\pi(A)$ (for any given $A \subseteq T_r \cap E_0$ with $|\pi(A) \setminus E_0| \leq 1$) in time $O(n)$, where n is the number of vertices of the graph G. Indeed, for $e = \{s, v\} \in A$, to be able to identify $\pi(e)$, one has to find the e-branch of T_r, and the v–s path in T_b. Both can be found by a depth-first search on the edge set $T_r \cup T_b$ of order $O(n)$. Different edges in A determine different branches of T_r, so every edge of T_r is visited at most once in a depth-first search. Moreover, because $|\pi(A) \setminus E_0| \leq 1$, at most two elements of $\pi(A)$ are in the same branch of T_b, so no branch of T_b has to be searched more than twice for a path.

Proof of Theorem 1. Necessity is dealt with in the introduction. To see sufficiency, suppose that both G_S and G_T contain k edge-disjoint spanning trees. Now we can apply Theorem 4 with E_1 the set of edges with at least one end in S, and E_2 the set of edges with at least one end in T. Indeed, by

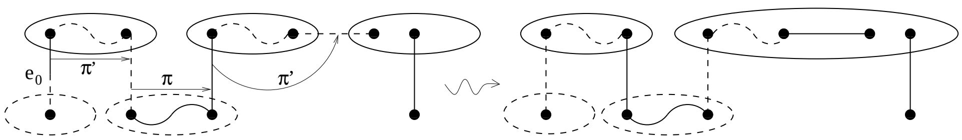

Figure 1: A step in the algorithm for packing connectors

Lemma 2, $E_0 = \delta(S)$ is an sbo-set for $\mathcal{M}(G_S)$ and for $\mathcal{M}(G_T)$, both viewed as matroids on $E = E_1 \cup E_2$ (observe that $E_0 = \delta(s)$ in G_S and $E_0 = \delta(t)$ in G_T). It follows that E contains k disjoint common spanning sets of $\mathcal{M}(G_S)$ and $\mathcal{M}(G_T)$, that is k edge-disjoint S–T connectors. □

In Figure 2, a typical step in the algorithm for packing connectors is shown. The vertices in S and T are the lower and upper vertices, respectively. Spanning trees N_r and N_b of G_S are indicated by dashed and solid lines (different 'colours'), respectively. So are the spanning trees M_r and M_b of G_T. Every S–T edge has two colours: the colour of the upper half depends on the spanning tree of G_S the edge belongs to, the colour of the lower half is determined by the spanning tree of G_T it belongs to. Branches of the trees in $G_S - s$ and $G_T - t$ are represented by ellipses of the corresponding colour in the picture. Paths consisting possibly of more than one edge are drawn wavy.

On the left, we have the situation that some S–T edge e_0 is in $N_b \cap M_r$. With the help of the sbo-injections π and π', as defined in Lemma 2, we find a sequence, as defined in Theorem 4, of edges alternatingly in $N_r \cap M_r$ and $M_b \cap N_b$, and in this case ending with an edge in $N_r \setminus M_r$. Swapping the colours of the edges in this sequence gives the situation on the right, where one more S–T edge, namely e_0, has both its halfs the same colour.

By the complexity analysis after Theorem 4 and after Lemma 2, we obtain an $O(\tau(n, m) + m'n)$ time algorithm for finding a maximum number of disjoint S–T connectors in the graph G on n vertices, m edges, where $m' := |\delta(S)| \leq m$. Here, $\tau(x, y)$ denotes the complexity of finding the maximum number of edge-disjoint spanning trees in a graph on x vertices and y edges. An algorithm for packing spanning trees due to Gabow and Westermann [2] proves that $\tau(n, m) \leq O(mn \log(m/n))$. So the bound we obtain for packing S–T connectors does not exceed this best known bound for packing spanning trees.

References

[1] J. Davies and C.J.H. McDiarmid, Disjoint common transversals and exchange structures, *J. London Math. Soc. (2)* **14** (1976), 55–62.

[2] H.N. Gabow and H.H. Westermann, Forests, frames, and games: algorithms for matroid sums and applications, *Algorithmica* **7** (1992), 465–497.

[3] R.P. Gupta, A decomposition theorem for bipartite graphs, in: *Theory of Graphs* (P. Rosenstiehl, ed.), Gordon and Breach, New York, 1967, 135–138.

[4] J. Keijsper and A. Schrijver, On packing connectors, *J. Combinatorial Theory Ser. B*, to appear.

[5] D. Kőnig, Űber Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre, *Math. Ann.* **77** (1916), 453–465.

[6] C.St.J.A. Nash-Williams, Edge-disjoint spanning trees of finite graphs, *J. London Math. Soc.* **36** (1961), 445–450.

[7] W.T. Tutte, On the problem of decomposing a graph into n connected factors, *J. London Math. Soc.* **36** (1961), 221–230.

[8] D.J.A. Welsh, *Matroid theory* , Academic Press, London, 1976.