COMBINATORICA
Bolyai Society – Springer-Verlag

COMBINATORICA **27** (1) (2007) 71–90

AN APPROXIMATE MAX-STEINER-TREE-PACKING MIN-STEINER-CUT THEOREM*

LAP CHI LAU[†]

Received April 6, 2004

Given an undirected multigraph G and a subset of vertices $S \subseteq V(G)$, the Steiner Tree Packing problem is to find a largest collection of edge-disjoint trees that each connects S. This problem and its generalizations have attracted considerable attention from researchers in different areas because of their wide applicability. This problem was shown to be APX-hard (no polynomial time approximation scheme unless P = NP). In fact, prior to this paper, not even an approximation algorithm with asymptotic ratio $o(n)$ was known despite several attempts.

In this work, we present the first polynomial time constant factor approximation algorithm for the Steiner Tree Packing problem. The main theorem is an approximate min-max relation between the maximum number of edge-disjoint trees that each connects S (S-*trees*) and the minimum size of an edge-cut that disconnects some pair of vertices in S (S-*cut*). Specifically, we prove that if every S-cut in G has at least $26k$ edges, then G has at least k edge-disjoint S-trees; this answers Kriesell's conjecture affirmatively up to a constant multiple.

1. Introduction

We consider a well-studied generalization of the edge-disjoint a,b-paths problem, namely, the Steiner Tree Packing problem. Given an undirected multigraph $G=(V,E)$ and $S \subseteq V(G)$. We say the vertices in S are *black*

Mathematics Subject Classification (2000): 05C05, 05C40, 05C70, 68R10, 68W25

* A preliminary version appeared in the Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS) 2004.

[†] The author was supported by an Ontario Graduate Scholarship and a University of Toronto Fellowship.

(also known as *terminal* vertices) while the vertices in $V(G)-S$ are *white* (also known as *Steiner* vertices); an edge is *white* if it connects two white vertices. A *S-Steiner-tree* (*S-tree*) is a tree of G that contains every vertex in S, a *S-Steiner-cut* (*S-cut*) is a subset of edges whose removal disconnects some vertices in S. The Steiner Tree Packing problem is to find a largest collection of edge-disjoint S-trees of G.

This problem and its generalization (where different specified subsets of vertices have to be connected by edge-disjoint trees) have attracted considerable attention from researchers in different areas. The Steiner Tree Packing problem has applications in routing problems arising in VLSI circuit design [19, 26, 32, 10–14, 37, 16], where an effective way of sharing different signals amongst cells in a circuit can be achieved by the use of edge-disjoint Steiner trees. It also has a variety of computer network applications such as multicasting [29, 2–4, 38, 9], video-conferencing [15] and network information flow [34, 24], where simultaneous communications can be facilitated by using edge-disjoint Steiner trees.

When $S=V(G)$, the Steiner Tree Packing problem is known as the Spanning Tree Packing problem. Tutte [36] and Nash-Williams [28] independently proved that a graph has k edge-disjoint spanning trees if and only if $E_G(\mathcal{P}) \geq k(|\mathcal{P}|-1)$ for every partition $\mathcal{P}$ of $V(G)$ into nonempty classes, where $E_G(\mathcal{P})$ denotes the number of edges connecting distinct classes of $\mathcal{P}$. As a corollary of Tutte and Nash-Williams result, every $2k$-edge-connected graph has k edge-disjoint spanning trees. Karger [18] exploited this approximate min-max relation to give the best known algorithm (near linear time) to compute a minimum cut of a graph. In fact, the Spanning Tree Packing problem is best investigated within the structures offered by matroids (see [35]), where Edmonds' matroid partition theorem yields a short proof of Tutte and Nash-Williams theorem as a corollary.

The Steiner Tree Packing problem, however, is NP-complete. Therefore, under the assumption that NP $\neq$ co-NP, a min-max relation like the Tutte–Nash-Williams theorem does not exist. Nonetheless, Kriesell [20, 21] conjectures that the approximate min-max corollary of the Tutte–Nash-Williams theorem does generalize to the Steiner Tree Packing problem. In the following, we say a graph is *k-S-connected* if every S-cut has at least k edges.

Kriesell's conjecture ([20, 21]). *If G is $2k$-S-connected, then G has k edge-disjoint S-trees.*

The conjecture is best possible for every k as shown by any k-regular k-edge-connected graph G with $S=V(G)$.

1.1. Previous Work

Prior to this work, Kriesell's conjecture was wide open despite several attempts. It was not known to be true even when $2k$ is replaced by any $o(n)\cdot k$ (not even when $k = 2$ [17]). Similarly, not even a polynomial time $o(n)$ approximation algorithm was known for the Steiner Tree Packing problem. That is, in simple graphs, no known polynomial time approximation algorithm has an asymptotic performance better than the naive algorithm of simply finding one spanning tree.

In the special case where every white vertex has an even degree, Kriesell [21] proves that his conjecture is true. An interesting corollary of this result is: if G is $2k$-S-connected, there is a collection of $2k$ S-trees such that every edge is used by at most 2 such S-trees; in other words, we have a 2-approximation algorithm if we allow *half integral solutions*. Also, the special case where there are no white edges is considered by Kriesell [20] and Frank, Király and Kriesell [8]. In particular, it is proven in [20] that if G has no white edge and G is $(k+1)k$-S-connected, then G has k edge-disjoint S-trees. This result is improved in [8] by replacing $(k+1)k$ with $3k$; it is based on a generalization of the Tutte–Nash-Williams theorem to hypergraphs using matroid theory. Recently, Kriesell [22] proves that if G is $(l+2)k$-S-connected where l is the maximum size of a *bridge* (see [22] for the definition), then G has k edge-disjoint S-trees; this result is a common generalization of the Tutte–Nash-Williams theorem (when $l=0$) and the case where there is no white edge (when $l=1$).

For the general case, Petingi and Rodriguez [31] prove that if G is $2(\frac{3}{2})^{|V(G)-S|}k$-$S$-connected, then G has k edge-disjoint S-trees. Kriesell [21], by using the result for the case that every white vertex has an even degree, improves this by weakening the connectivity requirement to $2|V(G)-S|+2k$. Jain, Mahdian and Salavatipour [17], by using a *shortcutting* procedure, prove that if G is $(|S|/4+o(|S|))k$-S-connected, then G has k edge-disjoint S-trees; this improves an exponential connectivity bound in terms of $|S|$ obtained earlier by Kriesell [21]. In both papers [21, 17], an optimal bound of $\lceil\frac{4}{3}k\rceil$ on the connectivity requirement is obtained for the case $|S|=3$.

Jain, Mahdian, Salavatipour also study a natural linear programming relaxation of the Steiner Tree Packing problem. The Fractional Steiner Tree Packing problem is formulated [17] by the following linear program. In the following $\mathcal{T}$ denotes the collection of all S-trees in a graph G, and c_e is the given *capacity* of the edge e.

$$\text{(1)} \qquad \begin{array}{ll} \text{maximize} & \sum_{T\in\mathcal{T}} x_T \\ \text{subject to} & \forall e\in E: \sum_{T\in\mathcal{T}} x_T \le c_e \qquad \forall T\in\mathcal{T}: x_T \ge 0 \end{array}$$

By using the Ellipsoid algorithm on the dual of the above linear program, Jain, Mahdian and Salavatipour [17] show that there is a polynomial time α-approximation algorithm for the Fractional Steiner Tree Packing problem if and only if there is a polynomial time α-approximation algorithm for the Minimum Steiner Tree problem. The Minimum Steiner Tree problem is to find a minimum weight S-tree for a given weighted graph. Robins and Zelikovsky [33] give a 1.55 approximation algorithm, and Bern and Plassmann [1] show that it is APX-hard (no polynomial time approximation scheme unless P=NP). Therefore, by using the results of the Minimum Steiner Tree problem, the Fractional Steiner Tree Packing problem is APX-hard but can be approximated within a factor of 1.55 to the optimal solution [17]. As a consequence, the (integral) Steiner Tree Packing problem is shown to be APX-hard [17].

Besides designing approximation algorithms, effort has been put in to designing faster exact algorithms by integer programming approaches [26, 32, 10–14, 37, 16] as well as designing practical heuristic methods [29, 2–4, 38, 9, 15, 34].

1.2. Our Contributions

The major contribution of this paper is the following approximate max-S-tree-packing min-S-cut theorem, which answers Kriesell's conjecture affirmatively up to a constant multiple.

Theorem 1.1. *If G is $26k$-S-connected, then G has k edge-disjoint S-trees.*

The proof of Theorem 1.1 is based on a new idea of graph decomposition, the edge splitting lemma by Mader [25] and a result by Frank, Király and Kriesell [8]. The proof is constructive so if G is $26k$-S-connected, then a collection of k edge-disjoint S-trees can be constructed in polynomial time. This implies the first polynomial time constant factor approximation algorithm for the Steiner Tree Packing problem. In the following, $\lambda_S(G)$ denotes the size of a minimum S-cut in G.

Theorem 1.2. *There is a polynomial time algorithm to construct a collection of at least $\lfloor \frac{\lambda_S(G)}{26} \rfloor$ edge-disjoint S-trees.*

The Capacitated Steiner Tree Packing problem is a generalization of the Steiner Tree Packing problem where each edge e has an integer capacity c_e which bounds the number of trees that can use e (the Steiner Tree Packing problem is the special case where $c_e = 1$ for all

$e \in E(G)$). Notice that LP(1) is a relaxation of the Capacitated Steiner Tree Packing problem and the optimal fractional solution to LP(1) of G is bounded above by the minimum capacity of a S-cut. By replacing each edge e of G by c_e multiple edges and applying Theorem 1.1 on the resulting graph, say G', we give the first constant upper bound on the integrality gap of (1).

Corollary 1.3. *The integrality gap of LP(1) is bounded above by 51.*

Applying Theorem 1.2 on G', however, only gives a pseudo-polynomial time approximation algorithm for the Capacitated Steiner Tree Packing problem to G. Nonetheless, by combining the approximation algorithm for the Fractional Steiner Tree Packing problem in [17] and the algorithm of Theorem 1.2, we are able to obtain a polynomial time algorithm for the Capacitated Steiner Tree Packing problem which constructs an integral solution of value at least $\lfloor \frac{\tau}{26} \rfloor$ (see Section 4), where τ is the value of an optimal integral solution.

2. Overview and the Setup

To understand our approach, it is illuminating to start from the ground work. In [8], Frank, Király and Kriesell consider a hypergraph generalization of the Spanning Tree Packing problem. A hypergraph H is *k-partition-connected* if $E_H(\mathcal{P}) \geq k(|\mathcal{P}|-1)$ holds for every partition $\mathcal{P}$ of $V(H)$ into non-empty classes, where $E_H(\mathcal{P})$ denotes the number of hyperedges intersecting at least two classes. The main theorem in [8] states that a hypergraph is k-partition-connected if and only if H can be decomposed into k sub-hypergraphs each of which is 1-partition-connected. The proof is based on the observation that the hyperforests (see [8] for the definition of a hyperforest) of a hypergraph form the family of independent sets of a matroid and thus Edmonds' matroid partition theorem can be applied.

Now, suppose an instance of the Steiner Tree Packing problem where G has no white edge is given. We can assume every white vertex is of degree 3 in G by using Mader's splitting lemma (in Section 3.2). Now, we construct a hypergraph H with vertex set S. For every white vertex v, there is a corresponding hyperedge of size 3 in H consisting of the neighbours of v. Also, $uv \in E(H)$ if $u,v \in S$ and $uv \in E(G)$. By applying the min-max theorem on the hypergraph problem, the following result on the Steiner Tree Packing problem is obtained as a corollary.

Theorem 2.1 **([8]).** *If G has no white edge and is $3k$-S-connected, then G has k edge-disjoint S-trees.*

Given an instance of the Steiner Tree Packing problem, our method is to reduce the general case to the above seemingly restrictive case when there is no white edge. The key observation is that if Theorem 1.1 holds, it holds with a rich combinatorial property, which we call *the extension property*. The extension property roughly (formally defined in Section 2.1) says that for any edge-partition of the edges incident to a "small" degree vertex, the edge-partition can be extended to edge-disjoint S-trees such that each class in the edge-partition is contained in one S-tree.

The proof can be divided into two steps. Given a graph G with l white edges, we search for a minimum S-cut in G with a white edge, and decompose G through the cut, resulting in two graphs G_1 and G_2 with a total of at most $l-1$ white edges. The cut decomposition lemma (in Section 3.1) shows that if Theorem 1.1 holds in both G_1 and G_2 with the extension property, then we can always "piece together" the solutions in G_1 and G_2 so that Theorem 1.1 also holds in G with the extension property. Therefore, by applying the cut decomposition step recursively, we reduce an instance with l white edges to at most $l+1$ instances without a white edge. By the cut decomposition lemma, if all those $l+1$ graphs (without a white edge) satisfy Theorem 1.1 with the extension property, then G satisfies Theorem 1.1 (with the extension property) by "piecing" their solutions together. This key step removes the difficulty of having white edges, and gives new insight into the core of the problem. It should be mentioned that the Steiner Tree Packing problem remains APX-hard when there is no white edge.

The second step (in Section 3.3), of course, is to prove that Theorem 1.1 does indeed hold with the extension property when there are no white edges. By using Mader's splitting lemma, we can assume that every white vertex is of degree 3 (in Section 3.2); and this gives us a set of "good" paths. With a sufficiently high connectivity assumption ($26k$ in Theorem 1.1), by using Theorem 2.1, we show that the extension property holds for any graph without a white edge and with every white vertex of degree 3. This step is more technical, but the intuition is simple – when the graph is highly S-connected, we have much freedom to construct the edge-disjoint S-trees. And it turns out that any edge-partition of the edges incident to a "small" degree vertex can be extended to edge-disjoint S-trees. This completes the high level description of our approach.

2.1. The Setup

Let G be λ-S-connected, a *small vertex* is a black vertex of degree λ in G. Let $E(u)$ be the set of edges that are incident to u, $\mathcal{P}_k(u) = \{E_1, \ldots, E_k\}$

is a *balanced edge-subpartition* of u if $E_1 \cup E_2 \cup \ldots \cup E_k \subseteq E(u)$, $|E_i| \geq 2$ for $1 \leq i \leq k$, and $E_i \cap E_j = \emptyset$ for $i \neq j$. We denote the set of neighbours of u in E_i by $N_{E_i}(u)$. A subgraph H *spans* a subset of vertices U if $U \subseteq V(H)$. H is a S*-subgraph* of G if it is a connected subgraph of G that spans S, H is a *double* S*-subgraph* of G if it is a S-subgraph of G and every vertex in S is of degree at least 2 in H.

Definition 2.2 (The Extension Property). Given G, $S \subseteq V(G)$, and a balanced edge-subpartition $\mathcal{P}_k(v) = \{E_1, \ldots, E_k\}$ of a small vertex v. $\{H_1, \ldots, H_k\}$ are k edge-disjoint S-subgraphs that *extend* $\mathcal{P}_k(v)$ if for $1 \leq i \leq k$:

(1) $E_i \subseteq E(H_i)$;
(2) $H_i - v$ is a $(S-v)$-subgraph that spans $N_{E_i}(v)$.

Theorem 2.3 (The Extension Theorem). *If G is $26k$-S-connected, then G has k edge-disjoint double S-subgraphs. Furthermore, for any balanced edge-subpartition $\mathcal{P}_k(v)$ of any small vertex v, G has k edge-disjoint double S-subgraphs that extend $\mathcal{P}_k(v)$.*

It is clear that Theorem 2.3 implies Theorem 1.1 as we just need the first statement. Let $\mathcal{G}$, henceforth, be a counterexample to Theorem 2.3 with the minimum number of edges, and let $Q = 26$. Without loss of generality, we also assume that $\mathcal{G}$ is connected. Our plan, hence, is to show that $\mathcal{G}$ does not exist and thus Theorem 2.3 holds. The proof of Theorem 2.3 is divided into three parts. First, in Section 3.1, we prove that $\mathcal{G}$ has no white edge by using the cut decomposition lemma. Then, in Section 3.2, we prove that every white vertex of $\mathcal{G}$ is of degree 3 by using Mader's splitting lemma. Finally, in Section 3.3, we prove that the extension property does hold in $\mathcal{G}$ and thus $\mathcal{G}$ does not exist.

3. Proof of the Extension Property

3.1. Cut Decomposition

The following lemma is the key step mentioned previously, which reduces Theorem 2.3 from the general case to the case where there is no white edge. The cut decomposition operation will be described inside the proof.

Lemma 3.1 (The Cut Decomposition Lemma). *$\mathcal{G}$ has no white edge.*

Proof. Let e be a white edge. If $\mathcal{G} - e$ is still Qk-S-connected, then by the choice of $\mathcal{G}$, we get our desired edge-disjoint double S-subgraphs in $\mathcal{G} - e$ and thus in $\mathcal{G}$.

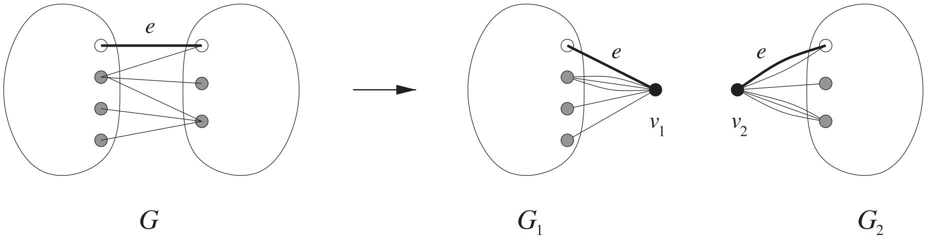

Figure 1. The construction of G_1 and G_2 from G.

Cut decomposition. So we consider the case that there is a S-cut $T=\{e_1,\ldots,e_{Qk}\}$ containing e. By the minimality of T, there are exactly two connected components C_1 and C_2 in $G-T$. Now we construct a new multigraph G_1 by contracting C_2 to a single black vertex v_1, keeping all edges from v_1 to C_1 (even if this produces multiple edges); similarly, we construct another new multigraph G_2 by contracting C_1 to a single black vertex v_2. So $V(G_1)=C_1+v_1$, $V(G_2)=C_2+v_2$, $T\subseteq E(G_1)$ and $T\subseteq E(G_2)$ (see Figure 1 for an illustration).

Let S_1 be the set of black vertices in C_1 plus v_1 and S_2 be the set of black vertices in C_2 plus v_2. Now we check the properties of G_1 and G_2. First, since e is in T, by contracting a component of size at least two (each component has at least one white vertex and one black vertex since e is white) to a single vertex, G_1 and G_2 have fewer edges than $\mathcal{G}$. Second, if $\mathcal{G}$ is Qk-S-connected, then G_1 is Qk-S_1-connected and G_2 is Qk-S_2-connected, since we keep multiple edges. Therefore, by the choice of $\mathcal{G}$, Theorem 2.3 holds in both G_1 and G_2. Note that v_1 and v_2 are small vertices since $|T|=Qk$, and G_1 and G_2 have a total of at most $l-1$ white edges if G has l white edges. This ends the description of the cut decomposition operation[1].

Let $v\in C_1$ be a small vertex of $\mathcal{G}$ and $\mathcal{P}_k(v)=\{E_1,\ldots,E_k\}$ be a balanced edge-subpartition of v. Our goal is to show that $\mathcal{G}$ has k edge-disjoint double S-subgraphs that extend $\mathcal{P}_k(v)$ (the case where G has no small vertex is similar but easier, we omit the details for brevity). And our plan is to combine k edge-disjoint double S_1-subgraphs in G_1 that extend $\mathcal{P}_k(v)$ and k edge-disjoint double S_2-subgraphs in G_2 that extend $\mathcal{R}_k(v_2)$ ($\mathcal{R}_k(v_2)$ to be defined) to obtain k edge-disjoint double S-subgraphs in $\mathcal{G}$ that extend $\mathcal{P}_k(v)$. Since Theorem 2.3 holds in G_1, we can find k edge-disjoint double S_1-subgraphs $\{H_1^1,\ldots,H_k^1\}$ of G_1 that extend $\mathcal{P}_k(v)$. Let F_i be the set of edges in H_i^1 that are incident to v_1. Since v_1 is a black vertex in G_1 and H_i^1 is a double S_1-subgraph, we have $|F_i|\geq 2$. Also, $F_i\cap F_j=\emptyset$ for $i\neq j$

[1] Similar constructions have been used in [28,25,30]

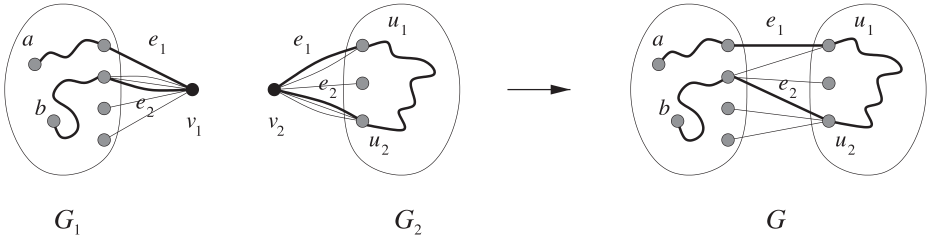

Figure 2. If $a,b\in C_1$ is connected in H_i^1 by a path through v_1 in G_1, they are still connected in H_i through C_2.

since H_i^1 and H_j^1 are edge-disjoint for $i\neq j$. Therefore, $\mathcal{R}_k(v_1)=\{F_1,\ldots,F_k\}$ is a balanced edge-subpartition of v_1 in G_1. Note that since v_1 and v_2 are incident to the same set of edges T, $\mathcal{R}_k(v_2)=\{F_1,\ldots,F_k\}$ is also a balanced edge-subpartition of v_2 in G_2. Since Theorem 2.3 holds in G_2, there are k edge-disjoint double S_2-subgraphs $\{H_1^2,\ldots,H_k^2\}$ of G_2 that extend $\mathcal{R}_k(v_2)$. We define a subgraph H_i of $\mathcal{G}$, by setting $E(H_i)$ to be the union of $E(H_i^1)$ and $E(H_i^2)$ with the exception that an edge of T in G_1 (or in G_2) becomes in H_i the corresponding edge in $\mathcal{G}$. We shall show that $\{H_1,\ldots,H_k\}$ are k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$.

First, notice that H_i^1 and H_i^2 use exactly the same edges in T, H_i^1 and H_j^1 are edge-disjoint for $i\neq j$, and H_i^2 and H_j^2 are edge-disjoint for $i\neq j$, so H_i and H_j are edge-disjoint for $i\neq j$. Now we shall show that H_i-v spans $N_{E_i}(v)$. Let $u\in N_{E_i}(v)$. If $u\in C_1$, u is spanned by H_i^1; if $u\in C_2$, then $u\in N_{F_i}(v_2)$ by our construction, so u is spanned by H_i^2. Therefore, H_i-v spans $N_{E_i}(v)$. Also, it follows from our construction that H_i-v spans $S-v$. So, to show that H_i–v is a $(S$–$v)$-subgraph of $\mathcal{G}$ that spans $N_{E_i}(v)$, it remains to show that H_i-v is a connected subgraph of $\mathcal{G}$. For any $a,b\in V(H_i)-v$, we consider the following three cases:

1. $a,b\in C_1$.
   If a and b are connected in H_i^1-v without using v_1, they are connected in H_i-v. So we consider the case that they are connected in H_i^1-v using v_1 (see Figure 2 for an illustration).
   Let e_1 and e_2 be the edges incident to v_1 in a path that connects a and b. Since $e_1,e_2\in E(H_i^1)\cap T$, by our construction, $e_1,e_2\in F_i$. Let u_1 and u_2 be the endpoints of e_1 and e_2 in C_2, so $u_1,u_2\in N_{F_i}(v_2)$. Recall that $H_i^2-v_2$ is a (S_2-v_2)-subgraph of G_2-v_2 that spans $N_{F_i}(v_2)$, so there is a path in $H_i^2-v_2$ between u_1 and u_2. By combining the edges in the a,v_1-path in H_i^1-v, the edges in the u_1,u_2-path in $H_i^2-v_2$ and the edges in the

v_1,b-path in $H_i^1 - v$, we get a path from a to b in $H_i - v$. As a result, a and b are connected in $H_i - v$.

2. $a \in C_1$, $b \in C_2$.
Since $H_i^1 - v$ is a $(S_1 - v)$-subgraph of $G_1 - v$, there is a a,v_1-path in $H_i^1 - v$. Let e be the edge incident to v_1 in the a,v_1-path. Since $e \in E(H_i^1) \cap T$, by our construction, $e \in F_i$. Let u be the endpoint of e in C_2. Since $H_i^2 - v_2$ is a $(S_2 - v_2)$-subgraph of $G_2 - v_2$ that spans $N_{F_i}(v_2)$, there is a u,b-path in $H_i^2 - v_2$. Therefore, there is a a,b-path in $H_i - v$ by combining the edges in the a,v_1-path and the edges in the u,b-path.
3. $a,b \in C_2$.
Recall that $H_i^2 - v_2$ is a $(S_2 - v_2)$-subgraph of $G_2 - v_2$, so a and b are connected in $H_i^2 - v_2$ and thus in $H_i - v$.

Therefore, $H_i - v$ is a $(S - v)$-subgraph that spans $N_{E_i}(v)$ (the second property of Definition 2.2 holds). By our construction, $E_i \subseteq E(H_i)$ (the first property of Definition 2.2 holds) which also implies that H_i is a S-subgraph of G. Furthermore, since u_1 is of degree at least 2 in H_i^1 for any $u_1 \in S_1$ and u_2 is of degree at least 2 in H_i^2 for any $u_2 \in S_2$, u is of degree at least 2 in H_i for any $u \in S$. Therefore, H_i is a double S-subgraph of $\mathcal{G}$. As a result, $\{H_1, \dots, H_k\}$ are k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$. Since v and $\mathcal{P}_k(v)$ are picked arbitrarily, this shows that Theorem 2.3 holds in $\mathcal{G}$, a contradiction. Therefore, $\mathcal{G}$ has no white edge and this completes the proof. ∎

3.2. Edge Splitting

A basic tool in the proof is Mader's splitting lemma, which is proven to be useful in many edge-connectivity problems. Let G be a graph, $e_1 = xy$, $e_2 = xz$ be two edges, $y \neq z$. The operation of obtaining $G(e_1, e_2)$ from G by deleting e_1 and e_2 and then adding exactly one new edge between y and z (multiple edges between y and z may be produced) is said to be *splitting at* x. This splitting at x is called *suitable*, if the number of edge-disjoint a,b-paths in $G(e_1, e_2)$ is at least the number of edge-disjoint a,b-paths in G for every pair $a,b \in V(G) - x$. Note that if we perform a suitable splitting at a white vertex, it does not decrease the S-connectivity. The splitting lemma provides a sufficient condition for the existence of a suitable splitting at a certain vertex x:

Lemma 3.2 (Mader's Splitting Lemma [25]). *Let x be a vertex of a graph G. Suppose that x is not a cut vertex and that x is incident with at least 4 edges and adjacent to at least 2 vertices. Then there exists a suitable splitting of G at x.*

Lemma 3.3. *There is no white cut vertex in $\mathcal{G}$.*

Proof. Suppose w is a white cut vertex in $\mathcal{G}$. Let $\{C_1,\ldots,C_l\}$ be the connected components of $\mathcal{G}-w$ where $l\geq 2$. We construct $G_i=\mathcal{G}[C_i\cup\{w\}]$ for $1\leq i\leq l$. Suppose all the black vertices are in only one component, say C_1. Since $\mathcal{G}$ is Qk-S-connected, G_1 is also Qk-S-connected and G_1 has fewer edges than G. So, by the choice of $\mathcal{G}$, Theorem 2.3 holds in G_1. But this implies that Theorem 2.3 also holds in $\mathcal{G}$, a contradiction.

So we assume that at least two components of $\mathcal{G}-w$ have black vertices. Let S_i be the black vertices in G_i. For any $a\in S_i$, since $\mathcal{G}$ is Qk-S-connected, it has Qk edge-disjoint paths to a vertex $b\in S_j$ for some $j\neq i$. Since w is a cut vertex, those Qk edge-disjoint a,b-paths must all pass through w. As a result, there are Qk edge-disjoint a,w-paths in G for any $a\in S_i$. This implies that each G_i is Qk-(S_i+w)-connected. By the choice of $\mathcal{G}$, each G_i has k edge-disjoint double $(S_i{+}w)$-subgraphs. By combining those k $(S_i{+}w)$-subgraphs of each G_i, we obtain k edge-disjoint double S-subgraphs of $\mathcal{G}$. Similarly, we can construct k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend any balanced edge-subpartition $\mathcal{P}_k(v)$ of any small vertex v (if any); a contradiction. Therefore, by the choice of $\mathcal{G}$, $\mathcal{G}$ has no white cut vertex. ∎

Lemma 3.4. *Every white vertex in $\mathcal{G}$ is incident with exactly three edges and adjacent to exactly three vertices.*

Proof. Suppose a white vertex w is adjacent to only one vertex u. Since $\mathcal{G}$ is Qk-S-connected, $\mathcal{G}-w$ is still Qk-S-connected. By the choice of $\mathcal{G}$, Theorem 2.3 holds in $\mathcal{G}-w$. Since u is not a small vertex, Theorem 2.3 also holds in $\mathcal{G}$, a contradiction. So we can assume that w is adjacent to at least two vertices.

Suppose a white vertex w is incident with only two edges, by the previous argument, w is adjacent to two vertices $\{y,z\}$. Since $\mathcal{G}$ is Qk-S-connected and $w\notin S$, $\mathcal{G}-w+yz$ is Qk-S-connected and it has one fewer edge than $\mathcal{G}$. By the choice of $\mathcal{G}$, Theorem 2.3 holds in $\mathcal{G}-w+yz$. For any k edge-disjoint double S-subgraphs $\{H_1,\ldots,H_k\}$ of $\mathcal{G}{-}w{+}yz$, if yz is in H_i, we can construct H_i' from H_i by replacing yz with $\{wy,wz\}$ so that H_i' is a double S-subgraph of $\mathcal{G}$. Note the remaining double S-subgraphs in $\mathcal{G}-w+yz$ are also double S-subgraphs in $\mathcal{G}$. So $\mathcal{G}$ has k edge-disjoint double S-subgraphs. Similarly, if the extension property holds in $\mathcal{G}-w+yz$, then the extension property holds in $\mathcal{G}$. But this implies that Theorem 2.3 holds in $\mathcal{G}$, a contradiction. So we can further assume that w is incident with more than two edges.

Suppose a white vertex w is incident with at least four edges. By the previous argument, w is adjacent to at least two vertices. And by Lemma 3.3, w is not a cut vertex. Therefore, by Lemma 3.2, there exists a suitable splitting of $\mathcal{G}$ at w, say the resulting graph is G^*. Since $\mathcal{G}$ is Qk-S-connected and the splitting is suitable, G^* is Qk-S-connected and has one fewer edge than $\mathcal{G}$. By the choice of $\mathcal{G}$, Theorem 2.3 holds in G^*. By a similar argument as in the previous paragraph, it follows that Theorem 2.3 also holds in $\mathcal{G}$; a contradiction. Therefore, the only possibility left is when w is incident with exactly three edges.

Suppose w is incident with three edges but adjacent to only two vertices $\{y,z\}$ so that there are two edges e_1,e_2 between w and y. Since $\mathcal{G}$ is Qk-S-connected, $w \notin S$ and w is incident with exactly three edges and adjacent only to $\{y,z\}$, it follows that $\mathcal{G}-e_1$ is Qk-S-connected and both w and y are not small vertices. By the choice of $\mathcal{G}$, Theorem 2.3 holds in $\mathcal{G}-e_1$. Since w and y are not small vertices, Theorem 2.3 also holds in $\mathcal{G}$, a contradiction. As a result, every white vertex w of $\mathcal{G}$ must be incident with exactly 3 edges and adjacent to exactly 3 vertices; this completes the proof. ∎

3.3. The Extension Property

Now we are ready to prove Theorem 2.3. The case when $|S|=2$ follows from Menger's theorem. Henceforth, we assume that $|S| \geq 3$. Let v be a small vertex, and $\mathcal{P}_k(v)=\{E_1,\ldots,E_k\}$ be a balanced edge-subpartition of v. Our goal, hence, is to show that $\mathcal{G}$ has k edge-disjoint double S-subgraphs that extend $\mathcal{P}_k(v)$. Let $W=\{w_1,\ldots,w_\alpha\}$ be the set of white neighbours of v and $B=\{b_1,\ldots,b_\gamma\}$ be the set of black neighbours of v. By Lemma 3.4, each w_i is incident with exactly three edges and adjacent to exactly three vertices, so we let $N_{\mathcal{G}}(w_i)=\{v,x_i,y_i\}$ and call $\{x_i,y_i\}$ a *couple*. Since w_i is a white vertex, by Lemma 3.1, x_i and y_i are black vertices. For each black neighbour b_i of v, the *weight* of b_i, denoted by $c(b_i)$, is the number of multiple edges between v and b_i.

Consider a black vertex $u \neq v$. Since $\mathcal{G}$ is Qk-S-connected, by Menger's theorem, there are Qk edge-disjoint paths, denoted by $P(u) = \{P_1(u),\ldots,P_{Qk}(u)\}$, from u to v. Note that since v is a small vertex, each path in $P(u)$ uses exactly one edge in $E(v)$. We assume vw_i is in the path $P_i(u)$ for $1 \leq i \leq \alpha$. Since w_i is of degree 3 by Lemma 3.4, $P_i(u)$ contains exactly one of w_ix_i or w_iy_i, and $P_j(u)$ does not contain w_ix_i or w_iy_i for $j \neq i$.

Let G' be $\mathcal{G}-v-W$. Consider $P_i(u)$ induced in G', denoted by $P'_i(u)$ (see Figure 3 for an illustration).

Let $P'(u)=\{P'_1(u),\ldots,P'_{Qk}(u)\}$, notice that $P'(u)$ contains edge-disjoint paths in G'. For $1 \leq i \leq \alpha$, $P'_i(u)$ is a path from u to either x_i or y_i in G'. Also, for each black neighbour b_j of v, there are $c(b_j)$ edge-disjoint paths in $P'(u)$ from u to b_j in G'.

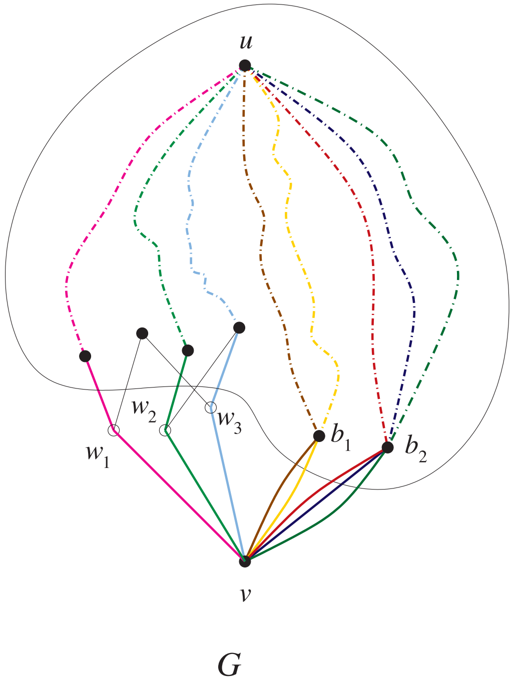

Figure 3. The paths in dotted lines are paths in $P'(u)$.

Let Z be a minimum $(S-v)$-cut of G' and $\{C_1,\ldots,C_l\}$ be the connected components of $G'-Z$. We let S_i and B_i be the set of black vertices and the set of black neighbours of v in C_i, respectively. Also, $c(B_i)$ denotes the sum of the weights of vertices in B_i and X_i denotes the collection of couples with both vertices in C_i. By the minimality of Z, each edge e in Z connects two vertices in different components, and we call it a *crossing edge*. Similarly, a couple $\{x_i,y_i\}$ is a *crossing couple* if x_i and y_i are in different components, and we denote the collection of crossing couples by X_C.

Now we give an outline of our proof of Theorem 2.3 when $\mathcal{G}$ has no white edge and every white vertex is of degree 3 and adjacent to exactly 3 vertices. We present the lemmas following the outline.

Outline. First, we show in Lemma 3.5 that if G' is $6k$-$(S-v)$-connected, then we can construct k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$ by using Theorem 2.1. Hence, by the choice of $\mathcal{G}$, we can assume G' has a $(S-v)$-cut Z so that $|Z|<6k$. Then, we show in Lemma 3.6 that $G'-Z$ has exactly 2 connected components C_1 and C_2, and in Lemma 3.7 that there are at least $Qk-2|Z|$ crossing couples. Consider any two black vertices $u_1,u_2\in C_i$, by using the paths in $P'(u_1)$ and $P'(u_2)$ and the above facts (i.e. Lemma 3.5 and Lemma 3.7), we show in Lemma 3.9 that there are at least $7k$ edge-disjoint paths from u_1 to u_2 in C_i. We further reserve at most k edges in each component to be used later. As a result, each component C_i is $6k$-S_i-connected and thus there are k edge-disjoint double S_i-subgraphs in C_i by Theorem 2.1. Finally, by exploiting the property that $\mathcal{P}_k(v)$ is a balanced edge-subpartition, we show in Lemma 3.10 that we can use the crossing edges in Z to connect the S_i-subgraphs to form k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$, a contradiction. This concludes the outline. ∎

Lemma 3.5. *G' is at most $(6k-1)$-$(S-v)$-connected.*

Proof. Since $|S|\geq 3$, $|S-v|\geq 2$. If G' is $6k$-$(S-v)$-connected, then there are $2k$ edge-disjoint $(S-v)$-subgraphs $\{H''_1,\ldots,H''_{2k}\}$ in G' by Theorem 2.1. Notice that since the union of two edge-disjoint $(S-v)$-subgraphs is a double $(S-v)$-subgraph, by setting $H'_i=H'_{2i-1}\cup H'_{2i}$, $\{H'_1,\ldots,H'_k\}$ are k edge-disjoint double $(S-v)$-subgraphs of G'. Now, let $H_i=H'_i\cup\{vb_j|vb_j\in E_i\}\cup\{vw_j,w_jx_j|vw_j\in E_i\}$. So, $E_i\subseteq H_i$, and H_i-v is a double $(S-v)$-subgraph that spans $N_{E_i}(v)$. Also, since H'_i is a double $(S-v)$-subgraph of G' and $|E_i|\geq 2$, H_i is a double S-subgraph of $\mathcal{G}$. By Definition 2.2, $\{H_1,\ldots,H_k\}$ are k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$, a contradiction. ∎

Lemma 3.6. *$G'-Z$ has 2 connected components.*

Proof. We just need to show that G' has at most 2 connected components, then the statement that $G'-Z$ has 2 connected components follows from the minimality of Z. Notice that from our construction of G' from $\mathcal{G}$, the set of neighbours of every white vertex that remained in G' is the same as in $\mathcal{G}$. Since $\mathcal{G}$ is connected, no component in G' contains only white vertices. Therefore, it suffices to show that there are at most two components in G' that contain black vertices.

Consider any two black vertices $u_1,u_2\neq v$. In $\mathcal{G}$, if v has a black neighbour b, then in G' there is a path in $P'(u_1)$ from u_1 to b and a path in $P'(u_2)$ from u_2 to b. So u_1 and u_2 are connected in G' and thus G' is connected. So suppose v has only white neighbours in $\mathcal{G}$. Consider $G''=G'+\{w_ix_i,w_iy_i\}$ for an arbitrary i, then the union of the edges in $P'_i(u_1)$, the edges in $P'_i(u_2)$ and $\{w_ix_i,w_iy_i\}$ contains a u_1,u_2-path in G''. Therefore, any two black vertices are in the same component in G'' and thus G'' is connected. Notice that w_i is a degree 2 vertex in G'', therefore $G'=G''-w_i$ has at most 2 connected components. ∎

Lemma 3.7. *There are at least $Qk-2|Z|$ crossing couples, that is, $|X_C|\geq Qk-2|Z|$.*

Proof. Let u_1 be a black vertex in C_1. In G', u_1 has at least $c(B_2)+|X_2|$ edge-disjoint paths in $P'(u_1)$ to C_2. Since Z is an edge-cut in G', it follows that $c(B_2)+|X_2|\leq|Z|$. Similarly, we have $c(B_1)+|X_1|\leq|Z|$. By Lemma 3.6, there are only two components in $G'-Z$. So $Qk=|X_C|+|X_1|+|X_2|+c(B_1)+c(B_2)$, and we have $|X_C|\geq Qk-2|Z|$. ∎

Now, we plan to use the paths in $P'(a)$ and $P'(b)$ for any two black vertices a,b in the same component of $G'-Z$ to establish the connectivity of each component of $G'-Z$. We say v_1 and v_2 have λ *common paths* if there are λ edge-disjoint paths starting from v_1, λ edge-disjoint paths starting from v_2, and an one-to-one mapping of the paths from v_1 to the paths from v_2 so that each pair of paths in the mapping ends in the same vertex. The following lemma gives a lower bound on the number of edge-disjoint paths between two vertices based on the number of their common paths, which will be used in Lemma 3.9 to prove that each C_i is $7k$-S_i-connected.

Lemma 3.8. *If v_1 and v_2 have $2\lambda+1$ common paths in G, then there exist $\lambda+1$ edge-disjoint paths from v_1 to v_2 in G.*

Proof. Suppose not, by Menger's theorem, there is an edge-cutset T of size at most λ that disconnects v_1 and v_2 in G. Since $|T|\leq\lambda$, at least $\lambda+1$ paths starting from v_1 remain in $G-T$; and the same holds for v_2. So, v_1 and v_2 have at least $(\lambda+1)+(\lambda+1)-(2\lambda+1)=1$ common path in $G-T$. This implies that v_1 and v_2 are connected in $G-T$, a contradiction. ∎

Lemma 3.9. *Each connected component C_i of $G'-Z$ is $7k$-S_i-connected.*

Proof. Let a,b be two black vertices in C_i. In G', $P'(a)$ has one path to each couple. Assume that, among those $|X_C|$ paths in $P'(a)$ to crossing couples, ϵ_a paths use edges in Z; and ϵ_b is defined similarly. Then, in $G'-Z$, a has $|X_C|-\epsilon_a$ edge-disjoint paths such that each starts from a and ends in a different crossing couple. Similarly, in $G'-Z$, b has $|X_C|-\epsilon_b$ edge-disjoint paths such that each starts from b and ends in a different crossing couple. Therefore, in $G'-Z$, a and b have at least $(|X_C|-\epsilon_a)+(|X_C|-\epsilon_b)-|X_C|=|X_C|-\epsilon_a-\epsilon_b$ pairs of paths that each pair of paths ends in the same crossing couple. Since a,b are in the same component, each such pair ends in the same endpoint of a crossing couple. So, a and b have at least $|X_C|-\epsilon_a-\epsilon_b$ common paths in C_i.

On the other hand, in G', $P'(a)$ has $c(B_2)+|X_2|$ edge-disjoint paths to C_2. Also, as mentioned in the previous paragraph, $P'(a)$ has ϵ_a edge-disjoint paths to crossing couples that use edges in Z. Notice that these $c(B_2)+|X_2|+\epsilon_a$ paths are edge-disjoint. Since Z is an edge-cut, Z has at least one edge in each such path. So, a has at least $c(B_2)+|X_2|+\epsilon_a$ edge-disjoint paths such that each path starts from a and ends in a different crossing edge in Z, note that they are also edge-disjoint from the paths mentioned in the previous paragraph. Similarly, $P'(b)$ has $c(B_2)+|X_2|+\epsilon_b$ edge-disjoint paths such that each path starts from b and ends in a different crossing edge in Z. Therefore, a and b have at least $(c(B_2)+|X_2|+\epsilon_a)+(c(B_2)+|X_2|+\epsilon_b)-|Z| = 2c(B_2)+2|X_2|+\epsilon_a+\epsilon_b-|Z|$ pairs of paths such that each pair of paths ends in the same crossing edge in Z. Since a and b are in the same component, each such pair of paths ends in the same endpoint of a crossing edge. So, a and b have at least $2c(B_2)+2|X_2|+\epsilon_a+\epsilon_b-|Z|$ more common paths in C_i.

As a result, by the previous two paragraphs, a and b have at least $2c(B_2)+ 2|X_2|+|X_C|-|Z|$ common paths in C_i. Recall that $c(B_2)+|X_2|+|X_C| = Qk-c(B_1)-|X_1|$ and $c(B_1)+|X_1|\leq|Z|$ (see the proof in Lemma 3.7), so a and b have at least $Qk+c(B_2)+|X_2|-2|Z|\geq Qk-2|Z|>(Q-12)k$ ($|Z|<6k$ by Lemma 3.5) common paths in C_i. Therefore, by Lemma 3.8, there are at least $(Q/2-6)k$ edge-disjoint a,b-paths in C_i. Since $Q=26$, this implies that C_i is $7k$-S_i-connected. ∎

Lemma 3.10. *$\mathcal{G}$ has k edge-disjoint double S-subgraphs $\{H_1,H_2,\dots,H_k\}$ that extend $\mathcal{P}_k(v)$.*

Proof. We pick arbitrarily $\min\{k,|Z|\}$ edges in Z and call them the *connecting edges*. For each connecting edge e with a white endpoint w in C_i, we remove one edge e' in C_i which is incident with w (by Lemma 3.1, the other endpoint of e' must be black), and we call e' a *reserve edge*. Let the resulting component be C'_i. Since we remove at most k edges and C_i is $7k$-S_i-connected by Lemma 3.9, each C'_i is $6k$-S_i-connected. By Theorem 2.1, there are $2k$ edge-disjoint S_i-subgraphs in C'_i. So there are k edge-disjoint double S_i-subgraphs $\{H^i_1,\dots,H^i_k\}$ in each C'_i for $i\in\{1,2\}$, except when $|S_i|=1$ for which we will consider separately later.

Now we set $H_j=H^1_j\cup H^2_j\cup\{vb_i|vb_i\in E_j\}\cup\{vw_i,w_ix_i,w_iy_i|vw_i\in E_j\}$ for $1\leq j\leq k$. Notice that $E_j\subseteq E(H_j)$ and H_j-v spans $N_{E_j}(v)$ for $1\leq j\leq k$. Suppose there is a crossing couple $\{x_i,y_i\}$ such that $vw_i\in E_j$, then H_j is also connected and thus is a S-subgraph of $\mathcal{G}$ that $E_j\subseteq E(H_j)$ and H_j-v is a $(S-v)$-subgraph that spans $N_{E_j}(v)$. Let's assume that $\{vw_1,\dots,vw_{|X_C|}\}$ be the set of edges such that the corresponding couples are crossing. By Lemma 3.7, $|X_C|\geq Qk-2|Z|$. Since $\mathcal{P}_k(v)$ is a balanced edge-subpartition, $|E_i|\geq 2$ for $1\leq i\leq k$. So, there are at most $\min\{k,|Z|\}$ classes of $\mathcal{P}_k(v)$ with no edges in $\{vw_1,\dots,vw_{Qk-2|Z|}\}$. Hence there are at most $\min\{k,|Z|\}$ of H_j's, say $\{H_1,\dots,H_{\min\{k,|Z|\}}\}$, are not connected by the crossing couples. Now, by adding each connecting edge and its reserve edge (if any) to a different H_j that has not been connected by a crossing couple, $\{H_1,\dots,H_k\}$ are k edge-disjoint S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$.

The only property left to be checked is if H_i is a double edge-disjoint S-subgraph for $1 \le i \le k$. Suppose $|S_1| \ge 2$, then every vertex $u \in S_1$ has degree at least 2 in every H_i since u has degree at least 2 in every H_i^1. The subtle case is $|S_1| = 1$, say $S_1 = \{x\}$, where each H_i^1 is trivial. Note that x is in every crossing couple in this case. Let $\{H_1,\dots,H_l\}$ be the S-subgraphs that x is a degree 1 vertex in them. Suppose $\{\{x,y_1\},\{x,y_2\},\dots,\{x,y_c\}\}$ are crossing couples such that $\{\{vw_1,xw_1,y_1w_1\},\dots,\{vw_c,xw_c,y_cw_c\}\} \subseteq E(H_j)$ and $\{vw_1,vw_2,\dots,vw_c\} \subseteq E_j$ and $c > 2$, then we can delete $\{xw_3,\dots,xw_c\}$ from H_j and do not affect the properties of H_j that are required in the preceding paragraph. We repeat this procedure until there are at least l edges, say $\{xw_1,\dots,xw_l\}$, that are not used in any H_j. Then we can add each such edge to a different S-subgraph in $\{H_1,\dots,H_l\}$ so that x is of degree at least 2 in each of $\{H_1,\dots,H_k\}$. We do the same "switching" procedure if $|S_2| = 1$. Since there are at least $Qk - 2|Z| > (Q-12)k = 14k$ crossing couples and there are only 2 components in $G' - Z$, the "switching" procedure is guaranteed to succeed. After all, $\{H_1,\dots,H_k\}$ are k edge-disjoint double S-subgraphs of $\mathcal{G}$ that extend $\mathcal{P}_k(v)$. ∎

Lemma 3.10 finishes the proof of Theorem 2.3 by showing that the minimum counterexample $\mathcal{G}$ does not exist.

4. Algorithmic Aspects and Generalization

The algorithm consists of two parts: The first step transforms the input graph G with l white edges to at most $l+1$ graphs $\{G_1,\dots,G_{l+1}\}$ such that each has no white edge, and every white vertex is of degree 3 and adjacent to exactly three black vertices. And the second step extends a balanced edge-subpartition of a small vertex in G_i to k edge-disjoint double S_i-subgraphs for each $1 \le i \le l+1$ and combines their solutions (where S_i is the set of black vertices in G_i). Theorem 2.1 can be solved by Edmonds' matroid partition algorithm [7,8]. The remaining steps can also be implemented in polynomial time, this justifies Theorem 1.2. Now, we use our algorithm and also the algorithm for the FRACTIONAL STEINER TREE PACKING problem to give a polynomial time approximation algorithm for the CAPACITATED STEINER TREE PACKING problem.

Theorem 4.1. *There is a polynomial time algorithm for the* CAPACITATED STEINER TREE PACKING *to construct an integral solution of value at least* $\lfloor \frac{\tau}{26} \rfloor$, *where* τ *is the value of an optimal integral solution.*

Proof. Given an instance of the Capacitated Steiner Tree Packing problem, let τ^*,τ be the value of an optimal fractional, integral solution, respectively. We first use the approximation algorithm for the Fractional Steiner Tree Packing problem [17] to obtain a fractional solution of value β such that $1.55\beta \geq \tau^*$. One feature of the above algorithm is that there are at most a polynomial number of trees in the fractional solution with $x_T > 0$, say $\{x_1,\dots,x_{p(n)}\}$. Suppose $\sum_{i=1}^{p(n)} \lfloor x_i \rfloor \geq \frac{1.55}{26}\sum_{i=1}^{p(n)} x_i$, then $\sum_{i=1}^{p(n)} \lfloor x_i \rfloor \geq \frac{1.55}{26}\sum_{i=1}^{p(n)} x_i = \frac{1.55}{26}\beta \geq \frac{1}{26}\tau^* \geq \frac{1}{26}\tau$. So, $\{\lfloor x_1 \rfloor,\dots,\lfloor x_{p(n)} \rfloor\}$ is an integral solution which is at least $\frac{\tau}{26}$, and we are done.

Otherwise, $\sum_{i=1}^{p(n)} x_i > \frac{26}{1.55}\sum_{i=1}^{p(n)} \lfloor x_i \rfloor$. Then, $(\frac{26}{1.55}-1)\sum_{i=1}^{p(n)} \lfloor x_i \rfloor < \sum_{i=1}^{p(n)} (x_i - \lfloor x_i \rfloor) \leq p(n)$, which implies $\sum_{i=1}^{p(n)} \lfloor x_i \rfloor < \frac{1.55}{26-1.55}p(n)$. So, $\beta = \sum_{i=1}^{p(n)} x_i = \sum_{i=1}^{p(n)} \lfloor x_i \rfloor + \sum_{i=1}^{p(n)} (x_i - \lfloor x_i \rfloor) < \frac{1.55}{26-1.55}p(n) + p(n) = \frac{26}{26-1.55}p(n)$. Therefore, $\tau^* < \frac{1.55\times 26}{26-1.55}p(n)$. Note that in any solution, at most a value of τ^* capacity is used in an edge (in other words, if $c_e > \tau^*$, the excess capacity $c_e - \tau^*$ will never be used). Now, to find an integral solution, we replace every edge e of G by $\min\{c_e, \lfloor \tau^* \rfloor\}$ multiple edges and call the resulting graph G'. Notice that the total number of edges in G' is bounded by a polynomial of n and the value of an optimal solution in G' is the same as in G. So, we can apply the algorithm in Theorem 1.2 to obtain $\lfloor \frac{\tau}{26} \rfloor$ edge-disjoint S-trees of G' in polynomial time, which correspond to an integral solution of G with value at least $\lfloor \frac{\tau}{26} \rfloor$. Therefore, in either case, the integral solution constructed is at least $\lfloor \frac{\tau}{26} \rfloor$. ∎

Acknowledgment

I am very grateful to Michael Molloy, Avner Magen, Derek Corneil, and Matthias Kriesell for invaluable comments; to Babak Farzad and Mohammad Salavatipour for fruitful discussions, and to Zongpeng Li for motivating the problem.

References

[1] M. Bern and P. Plassmann: The Steiner problem with edge lengths 1 and 2, *Information Processing Letters* **32** (1989), 171–176.

[2] D. Bienstock and A. Bley: Capacitated network design with multicast commodities, *Proceedings in the 8th International Conference on Telecommunication Systems*, 2000.

[3] S. Chen, O. Günlük and B. Yener: Optimal packing of group multicastings, *IEEE INFOCOMM'98*, 1998.

[4] S. Chen, O. Günlük and B. Yener: The multicast packing problem, *IEEE/ACM Transactions on Networking (TON)* **8** (2000), 311–318.

[5] R. Cole and R. Hariharan: A fast algorithm for computing Steiner edge connectivity, *Proceedings of the thirty-fifth ACM symposium on Theory of computing (STOC)*, pp. 167–176, 2003.

[6] J. Cheriyan and M. Salavatipour: Hardness and approximation results for packing Steiner trees problems, *Proceedings of ESA, Lecture Notes of Computer Science*, Vol. **3221**, pp. 180–191, 2004.

[7] J. Edmonds: Minimum partition of a matroid into independent sets, *J. Res. Nat. Bur. Standards Sect.* **869** (1965), 67–72.

[8] A. Frank, T. Király and M. Kriesell: On decomposing a hypergraph into k connected sub-hypergraphs, *Discrete Applied Mathematics* **131** (2003), 373–383.

[9] N. Garg, R. Khandekar, K. Kunal and V. Pandit: Bandwidth maximization in multicasting, *11th Annual European Symposium on Algorithms (ESA)*, 2003.

[10] M. Grötschel, A. Martin and R. Weismantel: Packing Steiner trees: further facets; *European J. Combin.* **17** (1996), 39–52.

[11] M. Grötschel, A. Martin and R. Weismantel: Packing Steiner trees: polyhedral investigations, *Mathematical Programming* **72** (1996), 101–123.

[12] M. Grötschel, A. Martin and R. Weismantel: Packing Steiner trees: a cutting plane algorithm and computational results; *Mathematical Programming* **72** (1996), 125–145.

[13] M. Grötschel, A. Martin and R. Weismantel: Packing Steiner trees: separation algorithms; *SIAM J. Discrete Math.* **9** (1996), 233–257.

[14] M. Grötschel, A. Martin and R. Weismantel: The Steiner tree packing problem in VLSI-design, *Mathematical Programming* **78** (1997), 265–281.

[15] M. Hosseini and N. D. Georganas: Design of a multi-sender 3D videoconferencing application over an end system multicast protocol, *Proceedings of the eleventh ACM international conference on Multimedia*, pp. 480–489, 2003.

[16] G. W. Jeong, K. Lee, S. Park and K. Park: A branch-and-price algorithm for the Steiner tree packing problem, *Computers and Operations Research* **29** (2002), 221–241.

[17] K. Jain, M. Mahdian and M. R. Salavatipour: Packing Steiner trees, *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pp. 266–274, 2003.

[18] D. Karger: Minimum cuts in near-linear time, *Journal of the ACM (JACM)* **47** (2000), 46–76.

[19] B. Korte, H. J. Prömel and A. Steger: Steiner trees in VLSI-layout, *Paths, Flows, and VLSI-layout*, Springer-Verlag, pp. 185–214, 1990.

[20] M. Kriesell: Local spanning trees in graphs and hypergraph decomposition with respect to edge-connectivity, *Technical Report* **257**, University of Hannover, 1999.

[21] M. Kriesell: Edge-disjoint trees containing some given vertices in a graph, *J. Combin. Theory, Series B* **88** (2003), 53–63.

[22] M. Kriesell: Disjoint Steiner trees in graphs without large bridges, manuscript, 2004.

[23] T. Leighton and S. Rao: Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, *Journal of the ACM (JACM)* **46** (1999), 787–832.

[24] Z. Li, B. Li, D. Jiang and L. C. Lau: On achieving optimal throughput with network coding, *IEEE INFOCOM'05*, 2005.
[25] W. Mader: A reduction method for edge-connectivity in graphs, *Ann. Discrete Math.* **3** (1978), 145–164.
[26] A. Martin and R. Weismantel: Packing paths and Steiner trees: routing of electronic circuits, *CWI Quarterly* **6** (1993), 185–204.
[27] K. Menger: Zur allgemeinen Kurventheorie, *Fund. Math.* **10** (1927), 95–115.
[28] C. St. J. A. Nash-Williams: Edge disjoint spanning trees of finite graphs, *J. London Math. Soc.* **36** (1961), 445–450.
[29] Y. Ofek and B. Yener: Reliable concurrent multicast from bursty sources, *IEEE Journal of Selected Areas in Communications* **15** (1997), 434–444.
[30] H. Okamura: Paths and edge-connectivity in graphs, *J. Combin. Theory, Series B* **37** (1984), 151–172.
[31] L. Petingi and J. Rodriguez: Bounds on the maximum number of edge-disjoint Steiner trees of a graph, *Congressus Numerantium* **145** (2000), 43–52.
[32] W. R. Pulleyblank: Two Steiner tree packing problems, *Proceedings of the twenty-seventh annual ACM Symposium on Theory of computing (STOC)*, pp. 383–387, 1995.
[33] G. Robins and A. Zelikovsky: Improved Steiner tree approximation in graphs, *Proceedings of the tenth annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 770–779, 2000.
[34] P. Sanders, S. Egner and L. Tolhuizen: Polynomial time algorithms for network information flow, *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*, pp. 286–294, 2003.
[35] A. Schrijver: *Combinatorial Optimization – Polyhedra and Efficiency*, Springer-Verlag, Berlin, 2003.
[36] W. T. Tutte: On the problem of decomposing a graph into n connected factors, *J. London Math. Soc.* **36** (1961), 221–230.
[37] E. Uchoa and M. Poggi de Aragão: Vertex-disjoint packing of two Steiner trees: polyhedra and branch-and-cut; *Mathematical Programming* **90** (2001), 537–557.
[38] C. F. Wang, C. T. Liang and R. H. Jan: Heuristic algorithms for packing of multiple-group multicasting, *Computers and Operations Research* **29** (2002), 905–924.

Lap Chi Lau

Department of Computer Science
University of Toronto
Sandford Fleming Building
10 King's College Road
Toronto, Ontario, M5S 3G4
Canada
chi@cs.toronto.edu