

On 2-Coverings and 2-Packings of Laminar Families

Joseph Cheriyan^{1*}, Tibor Jordán^{2**}, and R. Ravi^{3***}

¹ Department of Combinatorics and Optimization,
University of Waterloo, Waterloo ON Canada N2L 3G1,
e-mail: jcheriyan@math.uwaterloo.ca

² BRICS[†], Department of Computer Science, University of Aarhus,
Ny Munkegade, building 540, DK-8000 Aarhus C, Denmark.
e-mail: jordan@daimi.au.dk

³ GSIA, Carnegie Mellon University,
Pittsburgh, PA 15213-3890,
e-mail: ravi@cmu.edu

Abstract. Let \mathcal{H} be a laminar family of subsets of a groundset V . A k -cover of \mathcal{H} is a multiset C of edges on V such that for every subset S in \mathcal{H} , C has at least k edges that have exactly one end in S . A k -packing of \mathcal{H} is a multiset P of edges on V such that for every subset S in \mathcal{H} , P has at most $k \cdot u(S)$ edges that have exactly one end in S . Here, u assigns an integer capacity to each subset in \mathcal{H} .

Our main results are: (a) Given a k -cover C of \mathcal{H} , there is an efficient algorithm to find a 1-cover contained in C of size $\leq k|C|/(2k-1)$. For 2-covers, the factor of $2/3$ is best possible. (b) Given a 2-packing P of \mathcal{H} , there is an efficient algorithm to find a 1-packing contained in P of size $\geq |P|/3$. The factor of $1/3$ for 2-packings is best possible.

These results are based on efficient algorithms for finding appropriate colorings of the edges in a k -cover or a 2-packing, respectively, and they extend to the case where the edges have nonnegative weights. Our results imply approximation algorithms for some NP-hard problems in connectivity augmentation and related topics. In particular, we have a $4/3$ -approximation algorithm for the following problem: Given a tree T and a set of nontree edges E that forms a cycle on the leaves of T , find a minimum-size subset E' of E such that $T + E'$ is 2-edge connected.

1 Introduction

Let \mathcal{H} be a laminar family of subsets of a groundset V . In detail, let V be a groundset, and let $\mathcal{H} = \{S_1, S_2, \dots, S_q\}$ be a set of distinct subsets of V such

* Supported in part by NSERC research grant OGP0138432.

** Supported in part by the Hungarian Scientific Research Fund no. OTKA T29772 and T30059.

*** Supported in part by NSF career grant CCR-9625297.

[†] Basic Research in Computer Science, Centre of the Danish National Research Foundation.

that for every $1 \leq i, j \leq q$, $S_i \cap S_j$ is exactly one of \emptyset , S_i or S_j . A k -cover of \mathcal{H} is a multiset of edges, C , such that for every subset S in \mathcal{H} , C has at least k edges (counting multiplicities) that have exactly one end in S . A k -packing of \mathcal{H} is a multiset of edges, P , such that for every subset S in \mathcal{H} , P has at most $k \cdot u(S)$ edges (counting multiplicities) that have exactly one end in S . Here, u assigns an integer capacity to each subset in \mathcal{H} . Our main results are:

1. Given a k -cover C of \mathcal{H} , there is an efficient algorithm to find a 1-cover contained in C of size $\leq k|C|/(2k-1)$. For 2-covers, the factor of $2/3$ is best possible.
2. Given a 2-packing P of \mathcal{H} , there is an efficient algorithm to find a 1-packing contained in P of size $\geq |P|/3$. The factor of $1/3$ is best possible.

All of these results extend to the weighted case, where the edges have non-negative weights. Also, we show that the following two problems are NP-hard: (1) Given a 2-cover C of \mathcal{H} , find a minimum-size 1-cover that is contained in C . (2) Given a 2-packing P of \mathcal{H} , u , find a maximum-size 1-packing that is contained in P .

The upper bound of $2/3$ on the ratio of the minimum size of a 1-cover versus the size of a (containing) 2-cover is tight. To see this, consider the complete graph K_3 , and the laminar family \mathcal{H} consisting of three singleton sets. Let the 2-cover be $E(K_3)$. A minimum 1-cover has 2 edges from K_3 . The same example, with unit capacities for the three singleton sets in \mathcal{H} , shows that the ratio of the maximum size of a 1-packing versus the size of a (containing) 2-packing may equal $1/3$. There is an infinite family of similar examples.

An edge is said to *cover* a subset S of V if the edge has exactly one end in S . Our algorithm for finding a small-size 1-cover from a given 2-cover constructs a “good” 3-coloring of (the edges of) the 2-cover. In detail, the 3-coloring is such that for every subset S in the laminar family, at least two different colors appear among the edges covering S . The desired 1-cover is obtained by picking the two smallest (least weight) color classes. Similarly, our algorithm for finding a large-size 1-packing from a given 2-packing constructs a 3-coloring of (the edges of) the 2-packing such that for every subset S in the laminar family, at most $u(S)$ of the edges covering S have the same color. The desired 1-packing is obtained by picking the largest (most weight) color class.

1.1 A Linear Programming Relaxation

Consider the natural integer programming formulation (IP) of our minimum 1-cover problem. Let the given k -cover be denoted by E . There is a (nonnegative) integer variable x_e for each edge $e \in E$. For each subset $S \in \mathcal{H}$, there is a constraint $\sum_{e \in \delta(S)} x_e \geq 1$, where $\delta(S)$ denotes the set of edges covering S . The

objective function is to minimize $\sum_e w_e x_e$, where w_e is the weight of edge e . Let

(LP) be the following linear program obtained by relaxing all of the integrality

constraints on the variables.

$$(LP) \quad z_{LP} = \min \sum_e w_e x_e \quad \text{s.t.} \quad \left\{ \sum_{e \in \delta(S)} x_e \geq 1, \forall S \in \mathcal{H}; \quad x_e \geq 0, \forall e \in E \right\}.$$

Clearly, (LP) is solvable in polynomial time. The k -cover gives a feasible solution to (LP) by fixing $x_e = 1/k$ for each edge e in the k -cover.

For the minimum 1-cover problem, Theorem 1 below shows that the optimal value of the integer program (IP) is $\leq 4/3$ times the optimal value of a half-integral solution to the LP relaxation (LP). (A feasible solution x to (LP) is called *half-integral* if $x_e \in \{0, \frac{1}{2}, 1\}$, for all edges e .) There are examples where the LP relaxation has a unique optimal solution that is *not* half-integral. For the maximum 1-packing problem, Theorem 2 shows that the optimal value of the integer program is $\geq 2/3$ times the optimal value of a half-integral solution to the LP relaxation.

Recall that a laminar family \mathcal{H} may be represented as a tree $T = T(\mathcal{H})$. (T has a node for V as well as for each set $A_i \in \mathcal{H}$, and T has an edge $A_i A_j$ if $A_j \in \{V\} \cup \mathcal{H}$ is the smallest set containing $A_i \in \mathcal{H}$.)

Two special cases of the minimum 1-cover problem are worth mentioning.

- (i) If the laminar family \mathcal{H} is such that the tree $T(\mathcal{H})$ is a path, then the LP relaxation has an integral optimal solution. This follows because the constraints matrix of the LP relaxation is essentially a network matrix, see [CCPS 98, Theorem 6.28], and hence the matrix is totally unimodular; consequently, every extreme point solution (basic feasible solution) of the LP relaxation is integral.
- (ii) If the laminar family \mathcal{H} is such that the tree $T(\mathcal{H})$ is a star (i.e., the tree has one nonleaf node, and that is adjacent to all the leaf nodes) then the LP relaxation has a half-integral optimal solution. This follows because in this case the LP relaxation is essentially the same as the linear program of the fractional matching polytope, which has half-integral extreme point solutions, see [CCPS 98, Theorem 6.13].

1.2 Equivalent Problems

The problem of finding a minimum 1-cover of a laminar family \mathcal{H} from among the multiedges of a k -cover E may be reformulated as a connectivity augmentation problem. Let $T = T(\mathcal{H})$ be the tree representing \mathcal{H} ; note that $E(T)$ is disjoint from E . Then the problem is to find a minimum weight subset of edges E' contained in E such that $T + E' = (V(T), E(T) \cup E')$ is 2-edge connected; we may assume that E' has no multiedges. Instead of taking T to be a tree, we may take T to be a connected graph. This gives the problem *CBRA* which was initially studied by Eswaran & Tarjan [ET 76], and by Frederickson & Ja'ja' [FJ 81].

Similarly, the problem of finding a maximum 1-packing of a capacitated laminar family \mathcal{H} , u from among the multiedges of a k -packing E may be reformulated as follows. Let $T = T(\mathcal{H})$ be the tree representing \mathcal{H} , and let the tree edges have (nonnegative) integer capacities $u : E(T) \rightarrow \mathbf{Z}$; the capacity of a set

$A_i \in \mathcal{H}$ corresponds to the capacity of the tree edge a_i representing A_i . The k -packing E corresponds to a set of demand edges. The problem is to find a maximum integral multicommodity flow $x : E \rightarrow \mathbf{Z}$ where the source-sink pairs (of the commodities) are as specified by E . In more detail, the objective is to maximize the total flow $\sum_{e \in E} x_e$, subject to the capacity constraints, namely, for each tree edge a_i the sum of the x -values over the demand edges in the cut given by $T - a_i$ is $\leq u(a_i)$, and the constraints that x is integral and ≥ 0 .

1.3 Approximation Algorithms for NP-hard Problems in Connectivity Augmentation

Our results on 2-covers and 2-packings imply improved approximation algorithms for some NP-hard problems in connectivity augmentation and related topics. Frederickson and Ja'ja' [FJ 81] showed that problem *CBRA* is NP-hard and gave a 2-approximation algorithm. Later, Khuller and Vishkin [KV 94] gave another 2-approximation algorithm for a generalization, namely, find a minimum-weight k -edge connected spanning subgraph of a given weighted graph. Subsequently, Garg et al [GVY 97, Theorem 4.2] showed that problem *CBRA* is max SNP-hard, implying that there is no polynomial-time approximation scheme for *CBRA* modulo the $P \neq NP$ conjecture. Currently, the best approximation guarantee known for *CBRA* is 2.

Our work is partly motivated by the question of whether or not the approximation guarantee for problem *CBRA* can be improved to be strictly less than 2 (i.e., to $2 - \epsilon$ for a constant $\epsilon > 0$). We give a $4/3$ -approximation algorithm for an NP-hard problem that is a special case of *CBRA*, namely, the tree plus cycle (*TPC*) problem. See Section 4.

Garg, Vazirani and Yannakakis [GVY 97] show that the above maximum 1-packing problem (equivalently, the above multicommodity flow problem) is NP-hard and they give a 2-approximation algorithm. In fact, they show that the optimal value of an integral 1-packing z_{IP} is $\geq 1/2$ times the optimal value of a fractional 1-packing z_{LP} . We do not know whether the factor $1/2$ here is tight.

It should be noted that the maximum 1-packing problem for the special case of unit capacities (i.e., $u(A_i) = 1, \forall A_i \in \mathcal{H}$) is polynomial-time solvable. If the capacities are either one or two, and the tree $T(\mathcal{H})$ representing the laminar family \mathcal{H} has height two (i.e., every tree path has length ≤ 4), then the problem may be NP-hard, see [GVY 97, Lemma 4.3].

Further discussion on related topics may be found in the survey papers by Frank [F 94], Hochbaum [Hoc 96], and Khuller [Kh 96]. Jain [J 98] has interesting recent results, including a 2-approximation algorithm for an important generalization of problem *CBRA*.

We close this section by introducing some notation. For a multigraph $G = (V, E)$ and a node set $S \subseteq V$, let $\delta_E(S)$ denote the multiset of edges in E that have exactly one end node in S , and let $d_E(S)$ denote $|\delta_E(S)|$; so $d_E(S)$ is the number of multiedges in the cut $(S, V - S)$.

2 Obtaining a 1-Cover from a k -Cover

This section has our main result on k -covers, namely, there exists a 1-cover whose size (or weight) is at most $k/(2k-1)$ times the size (or weight) of a given k -cover. The main step (Proposition 1) is to show that there exists a “good” $(2k-1)$ -coloring of any k -cover. We start with a preliminary lemma.

Lemma 1. *Let V be a set of nodes, and let \mathcal{H} be a laminar family on V . Let E be a minimal k -cover of \mathcal{H} . Then there exists a set $X \in \mathcal{H}$ such that $d_E(X) = k$ and no proper subset Y of X is in \mathcal{H} .*

Proof. Since E is minimal, there exists at least one set $X \in \mathcal{H}$ with $d_E(X) = k$. We call a node set $X \subseteq V$ a *tight set* if $d_E(X) = k$. Consider an inclusionwise minimal tight set X in \mathcal{H} . Suppose there exists a $Y \subset X$ such that $Y \in \mathcal{H}$. If each edge of E that covers Y also covers X , then we have $d_E(Y) = k$. But this contradicts our choice of X . Thus there exists an edge $xy \in E$ covering Y with $x, y \in X$. By the minimality of E , xy must cover a tight set $Z \in \mathcal{H}$. Since \mathcal{H} is a laminar family, Z must be a proper subset of X . This contradiction to our choice of X proves the lemma. \square

Proposition 1. *Let V be a set of nodes, and let \mathcal{H} be a laminar family on V . Let E be a minimal k -cover of \mathcal{H} . Then there is a $(2k-1)$ -coloring of (the edges in) E such that*

- (i) *each set $X \in \mathcal{H}$ is covered by edges of at least k different colors, and*
- (ii) *for every node v with $d_E(v) \leq k$, all of the edges incident to v have distinct colors.*

Proof. The proof is by induction on $|\mathcal{H}|$. For $|\mathcal{H}| = 1$ the result holds since there are k edges in E (since E is minimal) and these can be assigned different colors. (For $|\mathcal{H}| = 0$, $|E| = 0$ so the result holds. However, even if E is nonempty, it is easy to color the edges in an arbitrary order to achieve property (ii).)

Now, suppose that the result holds for laminar families of cardinality $\leq N$. Consider a laminar family \mathcal{H} of cardinality $N+1$, and let E be a minimal k -cover of \mathcal{H} . By Lemma 1, there exists a tight set $A \in \mathcal{H}$ (i.e., $d_E(A) = k$) such that no $Y \subset A$ is in \mathcal{H} . We contract the set A to one node v_A , and accordingly update the laminar family \mathcal{H} . Then we remove the singleton set $\{v_A\}$ from \mathcal{H} . Let the resulting laminar family be \mathcal{H}' , and note that it has cardinality N . Clearly, E is a k -cover of \mathcal{H}' . Let $E' \subseteq E$ be a minimal k -cover of \mathcal{H}' . By the induction hypothesis, E' has a $(2k-1)$ -coloring that satisfies properties (i) and (ii), i.e., E' has a good $(2k-1)$ -coloring.

If the node v_A is incident to $\geq k$ edges of E' , then note that E' with its $(2k-1)$ -coloring is good with respect to \mathcal{H} (i.e., properties (i) and (ii) hold for \mathcal{H} too). To see this, observe that $k \leq d_{E'}(v_A) \leq d_E(v_A) = k$, so $d_{E'}(v_A) = k$, hence, the k edges of E' incident to v_A get distinct colors by property (ii). Then, for the original node set V , the k edges of E' covering A get k different colors.

Now focus on the case when $d_{E'}(v_A) < k$. Clearly, each edge in $E - E'$ is incident to v_A , since each edge in E not incident to v_A covers some tight set that

is in both \mathcal{H} and \mathcal{H}' . We claim that the remaining edges of $E - E'$ incident to v_A can be colored and added to E' in such a way that E with its $(2k - 1)$ -coloring is good with respect to \mathcal{H} .

It is easy to assign colors to the edge (or edges) of $E - E'$ such that the k edges of E incident to v_A get different colors. The difficulty is that property (ii) has to be preserved, that is, we must not “create” nodes of degree $\leq k$ that are incident to two edges of the same color. It turns out that this extra condition is easily handled as follows. Let $e \in E - E'$ be an edge incident to v_A , and let $w \in V$ be the other end node of e . If w has degree $\leq k$ for the current subset of E , then e is incident to $\leq (2k - 2)$ other edges; since $(2k - 1)$ colors are available, we can assign e a color different from the colors of all the edges incident to e . Otherwise (w has degree $> k$ for the current subset of E), the other edges incident to w impose no coloring constraint on e , and we assign e a color different from the colors of the other edges incident to v_A ; this is easy since $d_E(v_A) = k$. \square

Theorem 1. *Let V be a node set, and let \mathcal{H} be a laminar family on V . Let E be a k -cover of \mathcal{H} , and let each edge $e \in E$ have a nonnegative weight $w(e)$. Then there is a 1-cover of \mathcal{H} , call it E' , such that $E' \subseteq E$ and $w(E') \leq k w(E)/(2k - 1)$. Moreover, there is an efficient algorithm that given E finds E' ; the running time is $O(\min(k|V|^2, k^2|V|))$.*

Proof. We construct a good $(2k - 1)$ -coloring of the k -cover E by applying Proposition 1 to a minimal k -cover $\tilde{E} \subseteq E$ and then “extending” the good $(2k - 1)$ -coloring of \tilde{E} to E . That is, we partition E into $(2k - 1)$ subsets such that each set X in \mathcal{H} is covered by edges from at least k of these subsets. We take E' to be the union of the cheapest k of the $(2k - 1)$ subsets. Clearly, the weight of E' is at most $k/(2k - 1)$ of the weight of E , and (by property (i) of Proposition 1) E' is a 1-cover of \mathcal{H} .

Consider the time complexity of the construction in Proposition 1. Let $n = |V|$; then note that $|\mathcal{H}| \leq 2n$ and $|E| \leq 2kn$. The construction is easy to implement in time $O(|\mathcal{H}| \cdot |E|) = O(kn^2)$. Also, for $k < n$, the time complexity can be improved to $O(k^2 \cdot |\mathcal{H}|) = O(k^2n)$. To see this, note that for each set $A \in \mathcal{H}$ we assign colors to at most k of the edges covering A after we contract A to v_A , and for each such edge e we examine at most $(2k - 2)$ edges incident to e . \square

3 Obtaining a 1-Packing from a 2-Packing

This section has our main result on 2-packings, namely, there exists a 1-packing whose size (or weight) is at least $1/3$ times the size (or weight) of a given 2-packing. First, we show that there is no loss of generality in assuming that the 2-packing forms an Eulerian multigraph. Then we give a 3-coloring for the edges of the 2-packing such that for each set S in the laminar family at most $u(S)$ edges covering S have the same color. We take the desired 1-packing to be the biggest color class.

Lemma 2. *Let V be a set of nodes, let \mathcal{H} be a laminar family on V , and let $u : \mathcal{H} \rightarrow \mathbb{Z}$ assign an integral capacity to each set in \mathcal{H} . Let E be a 2-packing of \mathcal{H} , i.e., for all sets $A_i \in \mathcal{H}$, $d_E(A_i) \leq 2u(A_i)$. If E is a maximal 2-packing, then the multigraph $G = (V, E)$ is Eulerian.*

Proof. If G is not Eulerian, then it has an even number (≥ 2) of nodes of odd degree. Let $A \in \{V\} \cup \mathcal{H}$ be an inclusionwise minimal set that contains ≥ 2 nodes of odd degree. For every proper subset S of A that is in \mathcal{H} and that contains an odd-degree node, note that $d_E(S)$ is odd, hence, this quantity is strictly less than the capacity $2u(S)$. Consequently, we can add an edge (or another copy of the edge) vw where v, w are odd-degree nodes in A to get $E \cup \{vw\}$ and this stays a 2-packing of \mathcal{H} , u . This contradicts our choice of E , since E is a maximal 2-packing. Consequently, G has no nodes of odd degree, i.e., G is Eulerian. \square

Proposition 2. *Given an Eulerian multigraph $G = (V, E)$, an arbitrary pairing \mathcal{P} of the edges such that for every edge-pair the two edges have a common end node, and a laminar family of node sets \mathcal{H} , there is a 3-coloring of E such that*

- (i) *for each cut $\delta_E(A_i)$, $A_i \in \mathcal{H}$, at most half of the edges have the same color, and*
- (ii) *for each edge-pair e, f in \mathcal{P} , the edges e and f have different colors.*

Proof. Let \mathcal{P} be a set of triples $[v, e, f]$, where e and f are paired edges incident to the node v . Note that an edge $e = vw$ may occur in two triples $[v, e, f]$ and $[w, e, g]$. W.l.o.g. assume that \mathcal{P} gives, for each node v , a pairing of all the edges incident to v . Then \mathcal{P} partitions E into one or more (edge disjoint) subgraphs Q_1, Q_2, \dots , where each subgraph Q_j is a connected Eulerian multigraph. To see this, focus on the Eulerian tour given by fixing the successor of any edge $e = vw$ to be the other edge in the triple $[w, e, f] \in \mathcal{P}$, assuming e is oriented from v to w ; each such Eulerian tour gives a subgraph Q_j .

If $\mathcal{H} = \emptyset$, then we color each subgraph Q_j with 3 colors such that no two edges in the same edge-pair in \mathcal{P} get the same color. This is easy: We traverse the Eulerian tour of Q_j given by \mathcal{P} , and alternately assign the colors red and blue to the edges in Q_j , and if necessary, we assign the color green to the last edge of Q_j .

Otherwise, we proceed by induction on the number of sets in \mathcal{H} . We take an inclusionwise minimal set $A \in \mathcal{H}$, shrink it to a single node v_A , and update $G = (V, E)$, \mathcal{H} and \mathcal{P} to $G' = (V', E')$, \mathcal{H}' and \mathcal{P}' . Here, $\mathcal{H}' = \mathcal{H} - \{A\}$, i.e., the singleton set $\{v_A\}$ is not kept in \mathcal{H}' . Also, we add new edge pairs to \mathcal{P}' to ensure that all edges incident to v_A are paired. For a node $v \notin A$, all its triples $[v, e, f] \in \mathcal{P}$ are retained in \mathcal{P}' . Consider the pairing of all the edges incident to v_A in G' . For each triple $[v, e, f] \in \mathcal{P}$ such that $v \in A$ and each of e, f has one end node in $V - A$ (so e, f are both incident to v_A in G'), we replace the triple by $[v_A, e, f]$. We arbitrarily pair up the remaining edges incident to v_A in G' .

By the induction hypothesis, there exists a good 3-coloring for G' , \mathcal{H}' , \mathcal{P}' . It remains to 3-color the edges with both ends in A . For this, we shrink the nodes in $V - A$ to a single node v_B , and update $G = (V, E)$, \mathcal{P} , \mathcal{H} , to $G'' =$

$(V'', E''), \mathcal{P}'', \mathcal{H}''$; note that \mathcal{H}'' is the empty family and so may be ignored. We also keep the 3-coloring of $\delta_{E'}(v_A) = \delta_{E''}(v_B)$. Our final goal is to extend this 3-coloring to a good 3-coloring of E'' respecting \mathcal{P}'' . We must check that this can always be done. Consider the differently-colored edge pairs incident to v_B . Consider any connected Eulerian subgraph Q_j containing one of these edge pairs e_1, e_2 ; the corresponding triple in \mathcal{P}'' is $[v_B, e_1, e_2]$. Let \tilde{Q}_j be a minimal walk of (the Eulerian tour of) Q_j starting with e_2 and ending with an edge f incident to v_B (possibly, $f = e_1$). The number of internal edges in \tilde{Q}_j is $\equiv 0$ or $1 \pmod{2}$, and the two terminal edges either have the same color or not. If the number of internal edges in \tilde{Q}_j is nonzero, then it is easy to assign one, two, or three colors to these edges such that every pair of consecutive edges gets two different colors. The remaining case is when \tilde{Q}_j has *no* internal edges, say, $\tilde{Q}_j = v_B, e_2, w, f, v_B$, where w is a node in A . Then edges e_2, f are paired via the common end-node w , i.e., the triple $[w, e_2, f]$ is present in both \mathcal{P}'' and \mathcal{P} . Then, by our construction of \mathcal{P}' from \mathcal{P} , the triple $[v_A, e_2, f]$ is in \mathcal{P}' , and so edges e_2 and f (which are paired in \mathcal{P}' and present in $\delta_{E'}(v_A) = \delta_{E''}(v_B)$) must get different colors. Hence, a good 3-coloring of $G', \mathcal{H}', \mathcal{P}'$ can always be extended to give a good 3-coloring of \tilde{Q}_j , and the construction may be repeated to give a good 3-coloring of Q_j .

Finally, note that E'' is partitioned by \mathcal{P}'' into several connected Eulerian subgraphs Q_1, Q_2, \dots , where some of these subgraphs contain edges of $\delta_{E''}(v_B)$ and others do not. Clearly, the good 3-coloring of $G', \mathcal{H}', \mathcal{P}'$ can always be extended to give a good 3-coloring of each of Q_1, Q_2, \dots , and thus we obtain a good 3-coloring of $G, \mathcal{H}, \mathcal{P}$. \square

Theorem 2. *Let V be a node set, let \mathcal{H} be a laminar family on V , and let $u : \mathcal{H} \rightarrow \mathbf{Z}$ assign an integer capacity to each set in \mathcal{H} . Let E be a 2-packing of \mathcal{H} , and let each edge $e \in E$ have a nonnegative weight $w(e)$. Then there is a 1-packing of \mathcal{H} , call it E' , such that $E' \subseteq E$ and $w(E') \geq w(E)/3$. Moreover, there is an efficient algorithm that given E finds E' ; the running time is $O(|V| \cdot |E|)$.*

Proof. If the multigraph (V, E) is not Eulerian, then we use the construction in Lemma 2 to add a set of edges to make the resulting multigraph Eulerian without violating the 2-packing constraints. We assign a weight of zero to each of the new edges. Let us continue to use E to denote the edge set of the resulting multigraph. We construct a good 3-coloring of the 2-packing E by applying Proposition 2. Let F be the most expensive of the three “color classes;” so, the weight of F , $w(F)$, is $\geq w(E)/3$. Note that F is a 1-packing of \mathcal{H} , u by property (i) in the proposition since for every set $A_i \in \mathcal{H}$, we have $d_F(A_i) \leq d_E(A_i)/2 \leq 2u(A_i)$. Finally, we discard any new edges in F (i.e., the edges added by the construction in Lemma 2) to get the desired 1-packing.

Consider the time complexity of the whole construction. It is easy to see that the construction in Proposition 2 for the minimal set $A \in \mathcal{H}$ takes linear time. This construction may have to be repeated $|\mathcal{H}| = O(|V|)$ times. Hence, the overall running time is $O(|V| \cdot |E|)$. \square

4 Applications to Connectivity Augmentation and Related Topics

This section applies our covering result (Theorem 1) to the design of approximation algorithms for some NP-hard problems in connectivity augmentation and related topics. The main application is to problem *CBRA*, which is stated below. Problem *CBRA* is equivalent to some other problems in this area, and so we immediately get some more applications.

Recall problem *CBRA*: given a connected graph $T = (V, F)$, and a set of “supply” edges E with nonnegative weights $w : E \rightarrow \mathbb{R}_+$, the goal is to find a minimum-weight subset E' of E such that $T + E' = (V, F \cup E')$ is 2-edge-connected. One application of Theorem 1 is to give a $4/3$ -approximation algorithm for the special case of *CBRA* when the LP relaxation has an optimal solution that is half-integral.

Theorem 3. *Given a half-integral solution to the LP relaxation of CBRA of weight z , there is an $O(|V|)$ -time algorithm to find an integral solution (i.e., a feasible solution of CBRA) whose weight is $\leq \frac{4}{3}z$.*

Proof. Problem *CBRA* may be restated as the problem of finding a minimum-weight 1-cover of a laminar family \mathcal{H} , where the 1-cover must be chosen from the set of supply edges E and each supply edge has a nonnegative weight. To specify \mathcal{H} , fix any node $r \in V$ to be the root of T , and focus on the cut edges of T , call them f_1, f_2, \dots . For each of these cut edges f_i , let A_i be the (node set of the) component of $T - f_i$ that does not contain r . We take $\mathcal{H} = \{A_1, A_2, \dots\}$.

Let $x : E \rightarrow \{0, \frac{1}{2}, 1\}$ be a half-integral solution to the LP relaxation of *CBRA*, and let $z = \sum_e w_e x_e$. Then x corresponds to a 2-cover C of \mathcal{H} , where C has zero, one or two copies of a supply edge e iff $x_e = 0, 1$, or 2 . By Theorem 1, C contains a 1-cover C' whose weight is $\leq 4z/3$, and moreover, C' can be computed in time $O(|V|)$. \square

We have sharper results for the following (NP-complete) special case of problem *CBRA*.

Tree Plus Cycle Problem (TPC):

INSTANCE: A tree $T = (W, F)$ whose set of leaf nodes is $V \subseteq W$, a “supply” cycle $Q = (V, E)$ on the leaves of T (i.e., $d_E(v) = 2, \forall v \in V$), and a positive integer N .

QUESTION: Is there a set of edges $E' \subseteq E$ with $|E'| \leq N$ such that $T + E' = (W, F \cup E')$ is 2-edge-connected?

Corollary 1. *There exists a $\frac{4}{3}$ -approximation algorithm for TPC. Moreover, there exists a feasible solution $E' \subseteq E(Q)$ of size $\leq 2|V(Q)|/3$.*

Proof. Consider the LP relaxation of problem *TPC*; it is easy to verify that an optimal solution is given by $x_e = 1/2$ for all supply edges $e \in E(Q)$. Now, the result follows directly from Theorem 3. \square

Now consider the following problem: given a $(2k - 1)$ -edge-connected graph $T = (V, F)$ and a set of “supply” edges E with nonnegative weights $w : E \rightarrow \mathbb{R}_+$, the goal is to find a minimum weight subset E' of E such that $G' = (V, F + E')$ is $(2k)$ -edge-connected. Since the edge connectivity of T is odd, this problem is equivalent to problem *CBRA* because all the $(2k - 1)$ -cuts (minimum cuts) of T can be represented by means of a laminar family. (This follows easily from the fact that the node sets of two minimum cuts do not cross in this case.)

5 NP-completeness Results

First, we show that problem *TPC* (tree plus cycle) is NP-complete. It is convenient to reformulate *TPC* in terms of a laminar family rather than a tree.

Laminar Family Plus Cycle Problem (LPC):

INSTANCE: A laminar family \mathcal{H} on a node set V , a cycle $Q = (V, E)$ on V , and a positive integer N . (Assume $\emptyset, V \notin \mathcal{H}$.)

QUESTION: Is there a 1-cover E' of \mathcal{H} such that $E' \subseteq E$ and $|E'| \leq N$?

We give a polynomial-time reduction from the 3-dimensional matching problem to problem *LPC*. Our reduction is based on the proof of [FJ 81, Theorem 2].

Theorem 4. *Problem LPC is NP-complete.*

Proof. It is easy to see that *LPC* is in NP. Given an instance of *3DM* (that is, three disjoint sets W, X, Y , of cardinality q each, and a set M of 3-edges (triples) $(w_i x_j y_k) \in W \times X \times Y$), construct a connected graph T as follows. First build a star with a “root” r and $3q$ leaves $\{w_1, \dots, w_q, x_1, \dots, x_q, y_1, \dots, y_q\}$ corresponding to the elements of $W \cup X \cup Y$. Then for each 3-edge $(w_i x_j y_k)$ of M add two nodes a_{ijk} and \bar{a}_{ijk} to T and add the edges $w_i a_{ijk}$, $w_i \bar{a}_{ijk}$. Now replace each of the $2q$ nodes corresponding to elements of X and Y by complete graphs (or arbitrary 2-edge-connected graphs) denoted by $X_1, \dots, X_q, Y_1, \dots, Y_q$ as follows. Each complete subgraph of this type has $d_M(x_j)8q$ ($d_M(y_k)8q$) nodes and is partitioned into $d_M(x_j)$ ($d_M(y_k)$) parts (so-called “lanes”) of size $8q$ each. (Here, $d_M(x_j)$ and $d_M(y_k)$ denote the number of 3-edges of M containing x_j , respectively, y_k .) The graph constructed is connected and has $2p + 2q$ “leaves” (that is, leaf 2-edge-connected components), where $p := |M|$.

The next step is to define the cycle Q . The nodes of Q are the nodes of the leaves of T . Hence, $|V(Q)| = p(16q + 2)$. First, we define p disjoint paths of Q such that each has $16q + 2$ nodes (so each of these paths has length $16q + 1$). Every 3-edge $(w_i x_j y_k)$ of M defines such a path as follows: take the $8q$ nodes (and edges connecting the consecutive ones) $l_1 l_2 \dots l_{8q}$ of a lane of X_1 in an arbitrary order, then take the edges $l_{8q} a_{ijk}$, $a_{ijk} \bar{a}_{ijk}$, $\bar{a}_{ijk} m_{8q}$ for some node m_{8q} of some lane of Y_k , in this order, and then take the other nodes of this lane m_{8q-1}, \dots, m_1 in an arbitrary order. The lanes are chosen in such a way that these paths are pairwise disjoint. This can be done, since the lanes are pairwise disjoint and each X_j (or Y_k) has $d_M(x_j)$ (or $d_M(y_k)$) lanes.

Now fix a cyclic ordering e_1, \dots, e_p of the 3-edges of M and complete the cycle Q by adding the missing p edges in such a way that the end of the path corresponding to $e_s = (w_i x_j y_k)$ (that is, a node m_1 of a lane in Y_k) is connected to the first node of the path corresponding to $e_{s+1} = w_{i'} x_{j'} y_{k'}$ (that is, to a node l_1 of a lane of $X_{j'}$) for $1 \leq s \leq p$. Note that each of these edges connects a complete subgraph X_j to a complete subgraph Y_k and all the edges of Q either connect different leaves of T or connect different nodes of some leaf of T . Furthermore, $V(Q)$ equals the union of the nodes of the leaves of T .

The last part of the reduction consists of defining a laminar family \mathcal{H} on $V(Q)$. We define \mathcal{H} by defining two disjoint subfamilies \mathcal{H}_1 and \mathcal{H}_2 . Let $\mathcal{H}_1 := \{S \cap V(Q) : d_T(S) = 1, r \notin S, S \subseteq V(T)\}$ contain intersections of $V(Q)$ and those minimum cuts of T which do not contain the root. It is easy to see that this family is laminar. \mathcal{H}_2 consists of $2p$ disjoint collections, each of them defined on the nodes of a lane of a complete subgraph of the form X_j or Y_k of T as follows. Let us fix such a subgraph, say X_1 . (The definition is similar for all the $2q$ subgraphs $X_1, \dots, X_q, Y_1, \dots, Y_q$.) Focus on a lane l_1, \dots, l_{8q} of X_1 , where the numbering follows the ordering of these nodes in Q . (Hence l_{8q} is connected to some leaf a_{ijk} and l_1 is connected to some Y_k .) This lane adds the following sets to \mathcal{H}_2 : the singletons l_1, \dots, l_{8q} , the sets of nodes of the intervals of Q with end-node pairs (l_{8q-1}, l_{8q-s}) ($2 \leq s \leq 4q$) and (l_{4q-r}, l_2) ($1 \leq r \leq 4q-3$). Each lane of every complete subgraph X_j, Y_k ($1 \leq j, k \leq q$) adds a similar collection to \mathcal{H}_2 . Clearly, every collection of this type is laminar, and the collections are defined on pairwise disjoint sets of nodes, where each of these sets is included in a minimal element of \mathcal{H}_1 . Therefore \mathcal{H} is a laminar family on $V(Q)$, where $\mathcal{H} := \mathcal{H}_1 \cup \mathcal{H}_2$. Note that each node of Q belongs to \mathcal{H} as a singleton set.

Observe the following important property, that follows from the structure of these collections and the fact that every node of Q belongs to \mathcal{H} . Let $E' \subseteq E(Q)$ be a 1-cover of \mathcal{H} . Then

(*) if the edge $l_{8q} l_{8q-1}$ (or similarly $l_1 l_2, m_{8q} m_{8q-1}, m_1 m_2$) for some lane in an X_j or Y_k is not in E' then $|E'| \geq |V(Q)|/2 + 2q - 1$.

It is easy to see that our reduction is polynomial. We claim that there exists a solution to the given instance of $3DM$ (that is, a set of q pairwise disjoint 3-edges of M) if and only if \mathcal{H} has a 1-cover of size at most $p + 8pq + q = |V(Q)|/2 + q$.

First observe that a set E' is a 1-cover if and only if $T + E'$ is 2-edge-connected and E' covers each member of \mathcal{H}_2 . Moreover, as it was verified in [FJ 81], there is a 3-dimensional matching if and only if there is set E^* of $p + q$ edges in Q^* for which $T^* + E^*$ is 2-edge-connected, where T^* and Q^* arise from T and Q , respectively, by contracting the complete subgraphs (that is, the sets of the form X_j, Y_k , which are 2-edge-connected) to singletons and deleting the edges connecting these complete subgraphs from the cycle.

Suppose that there exists a 3-dimensional matching $M' \subseteq M$. Then there exists a set E^* of size $p + q$ which makes T^* 2-edge-connected and it is easy to see that there exists a set E'' of independent edges in Q which covers \mathcal{H}_2 . Hence $|E''| = 16qp/2 = 8pq$. Now $E' := E^* \cup E''$ covers \mathcal{H} and $|E'| = 8pq + p + q$, as required.

The proof of the other direction (which relies on $(*)$) is omitted. \square

Corollary 2. *The following problem is NP-hard: given a 2-cover C of a laminar family \mathcal{H} , find a minimum-size 1-cover that is contained in C .*

Theorem 5. *The following problem is NP-hard: given a 2-packing P of a capacitated laminar family \mathcal{H} , u , find a maximum-size 1-packing that is contained in P .*

6 Conclusions

We suspect that our bounds on the ratios for 1-covers versus 2-covers and for 1-packings versus 2-packings hold in general.

1-COVER CONJECTURE: Consider the integer program for a minimum weight 1-cover of a laminar family and its LP relaxation (see Section 1). We conjecture that the ratio of the optimal values is at most $4/3$.

1-PACKING CONJECTURE: Consider the integer program for a maximum weight 1-packing of a capacitated laminar family and its LP relaxation (see Section 1). We conjecture that the ratio of the optimal values is at least $2/3$.

Another interesting question is to find sufficient conditions on the laminar family \mathcal{H} (or, on the tree $T(\mathcal{H})$ representing \mathcal{H}) such that the LP relaxation has $\frac{1}{k}$ -integral extreme point solutions. As noted in Section 1, the LP relaxation has integral extreme point solutions iff $T(\mathcal{H})$ is a path.

References

- [CCPS 98] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver, *Combinatorial Optimization*, John Wiley & Sons, New York, 1998.
- [ET 76] K. Eswaran and R.E. Tarjan, "Augmentation problems," *SIAM J. Computing* **5** (1976), 653–665.
- [F 94] A. Frank, "Connectivity augmentation problems in network design," in *Mathematical Programming: State of the Art 1994*, (Eds. J. R. Birge and K. G. Murty), The University of Michigan, Ann Arbor, MI, 1994, 34–63.
- [FJ 81] G.N. Frederickson and J. Ja'Ja', "Approximation algorithms for several graph augmentation problems," *SIAM J. Comput.* **10** (1981), 270–283.
- [GVY 97] N. Garg, M. Yannakakis, and V. Vazirani, "Primal-dual approximation algorithms for integral flow and multicut in trees," *Algorithmica* **18** (1997), 3–20.
- [Hoc 96] D. S. Hochbaum, "Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems," in *Approximation algorithms for NP-hard problems*, Ed. D. S. Hochbaum, PWS co., Boston, 1996.
- [J 98] K. Jain, "A factor 2 approximation algorithm for the generalized Steiner network problem," Proc. 39th IEEE FOCS, Palo Alto, CA, November 1998.
- [Kh 96] S. Khuller, "Approximation algorithms for finding highly connected subgraphs," in *Approximation algorithms for NP-hard problems*, Ed. D. S. Hochbaum, PWS publishing co., Boston, 1996.
- [KV 94] S. Khuller and U. Vishkin, "Biconnectivity approximations and graph carvings," *Journal of the ACM* **41** (1994), 214–235.