

The Matching Augmentation Problem: A $\frac{7}{4}$ -Approximation Algorithm

J.Cheriyān * J.Dippel † F.Grandoni ‡ A.Khan § V.V.Narayan ¶

December 4, 2018

Abstract

We present a $\frac{7}{4}$ approximation algorithm for the matching augmentation problem (MAP): given a multi-graph with edges of cost either zero or one such that the edges of cost zero form a matching, find a 2-edge connected spanning subgraph (2-ECSS) of minimum cost.

We first present a reduction of any given MAP instance to a collection of well-structured MAP instances such that the approximation guarantee is preserved. Then we present a $\frac{7}{4}$ approximation algorithm for a well-structured MAP instance. The algorithm starts with a min-cost 2-edge cover and then applies ear-augmentation steps. We analyze the cost of the ear-augmentations using an approach similar to the one proposed by Vempala & Vetta for the (unweighted) min-size 2-ECSS problem (“Factor 4/3 approximations for minimum 2-connected subgraphs,” APPROX 2000, LNCS 1913, pp.262–273).

Keywords: 2-edge connected graph, 2-edge covers, approximation algorithms, bridges, connectivity augmentation, forest augmentation problem, matching augmentation problem, network design.

1 Introduction

A basic goal in the area of *survivable network design* is to design real-world networks of low cost that provide connectivity between pre-specified pairs of nodes even after the failure of a few edges/nodes. Many of the problems in this area are NP-hard, and significant efforts have been devoted in the last few decades to the design of approximation algorithms, see [20].

One of the fundamental problems in the area is the minimum-cost 2-edge connected spanning subgraph problem (abbreviated as min-cost 2-ECSS): given a graph together with non-negative costs for the edges, find a 2-edge connected spanning subgraph (abbreviated as 2-ECSS) of minimum cost. This problem is closely related to the famous Traveling Salesman Problem (TSP), and some of the earliest papers in the area of approximation algorithms address the min-cost 2-ECSS problem [5, 6]. In the context of approximation algorithms, this research led to the discovery of algorithmic paradigms such as the *primal-dual method* [8, 20] and the *iterative rounding method* [10, 13], and led to dozens of publications. Under appropriate assumptions, these methods achieve an approximation guarantee of 2 for several key problems in survivable network design, including min-cost 2-ECSS. Unfortunately, these generic methods do not achieve approximation guarantees

*University of Waterloo, Waterloo, Canada

†University of Waterloo, Waterloo, Canada

‡IDSIA, USI-SUPSI, Lugano, Switzerland

§Technical University of Munich, Garching, Germany

¶McGill University, Montreal, Canada

below 2. Significant research efforts have been devoted to achieving approximation guarantees below 2 for specific problems in the area of survivable network design. For example, building on earlier work, an approximation guarantee of $\frac{4}{3}$ has been achieved for unweighted (*min-size*) 2-ECSS [17], where each edge of the input graph has cost one and the goal is to find a 2-ECSS with the minimum number of edges.

There is an important obstacle beyond unweighted problems, namely, the special case of min-cost 2-ECSS where the (input) edges have cost of zero or one, and the aim is to design an algorithm that achieves an approximation guarantee below 2. This problem is called the *Forest Augmentation Problem* (FAP). In more detail, we are given an undirected graph $G = (V, E_0 \cup E_1)$, where each edge in E_0 has cost zero and each edge in E_1 has cost one; the goal is to compute a 2-ECSS $H = (V, F)$ of minimum cost. We denote the cost of an edge e of G by $\text{cost}(e)$, and for a subgraph G' of G , $\text{cost}(G')$ denotes $\sum_{e \in E(G')} \text{cost}(e)$. Observe that $\text{cost}(H) = |F \cap E_1|$, so the goal is to augment E_0 to a 2-ECSS by adding the minimum number of edges from E_1 . Intuitively, the zero-edges define some existing network that we wish to augment (with edges of cost one) such that the augmented network is resilient to the failure of any one edge. Without loss of generality (w.l.o.g.) we may contract each of the 2-edge connected subgraphs formed by the zero-edges, and hence, we may assume that E_0 induces a forest: this motivates the name of the problem.

A key special case of FAP is the *Tree Augmentation Problem* (TAP), where the edges of cost zero form a spanning tree. Nagamochi [15] first obtained an approximation guarantee below 2 for TAP, and since then there have been several advances including recent work, see [1, 3, 4, 9, 12].

We focus on a different special case of FAP called the *matching augmentation problem* (MAP): given a multi-graph with edges of cost either zero or one such that the edges of cost zero form a matching, find a 2-ECSS of minimum cost. Note that loops are not allowed; multi-edges (parallel edges) are allowed. From the view-point of approximation algorithms, MAP is “orthogonal” to TAP in the sense that the forest of zero-cost edges has many connected components in MAP, whereas this forest has only one connected component in TAP. In our opinion, MAP (like TAP) is an important special case of FAP in the sense that none of the previous approaches (including approaches developed for TAP over two decades) give an approximation guarantee below 2 for MAP.

1.1 Previous literature & possible approaches for MAP

There is extensive literature on approximation algorithms for problems in survivable network design, and also on the minimum-cost 2-ECSS problem including its key special cases (including the unweighted problem, TAP, etc.). To the best of our knowledge, there is no previous publication on FAP or MAP, although the former is well known to the researchers working in this area.

Let us explain briefly why previous approaches do not help for obtaining an approximation guarantee below 2 for MAP. Let G denote the input graph, and let n denote $|V(G)|$. Let opt denote the optimal value, i.e., the minimum cost of a 2-ECSS of the given instance. Recall that the standard cut-covering LP relaxation of the min-cost 2-ECSS problem has a non-negative variable x_e for each edge e of G , and for each nonempty set of nodes S , $S \neq V$, there is a constraint requiring the sum of the x -values in the cut $(S, V - S)$ to be ≥ 2 ; the objective is to minimize $\sum_{e \in E} \text{cost}(e)x_e$.

The primal-dual method and the iterative rounding method are powerful and versatile methods for rounding LP relaxations, but in the context of FAP, these methods seem to be limited to proving approximation guarantees of at least 2.

Several intricate combinatorial methods that may also exploit lower-bounds from LP relaxations have been developed for approximation algorithms for unweighted 2-ECSS, e.g., the $\frac{4}{3}$ -approximation algorithm of [17]. For unweighted 2-ECSS, there is a key lower bound of n on opt (since any solution must have $\geq n$ edges, each of cost one). This no longer holds for MAP; indeed,

the analogous lower bound on opt is $\frac{1}{2}n$ for MAP. It can be seen that an α -approximation algorithm for unweighted 2-ECSS implies a $(3\alpha - 2)$ -approximation algorithm for MAP. (We sketch the reduction: let M denote the set of zero-cost edges in an instance of MAP; observe that $|M| \leq opt$; we subdivide (once) each edge in M , then we change all edge costs to one, then we apply the algorithm for unweighted 2-ECSS, and finally we undo the initial transformation; the optimal cost of the unweighted 2-ECSS instance is $\leq opt + 2|M|$, hence, the solution of the MAP instance has cost $\leq \alpha(opt + 2|M|) - 2|M| = \alpha opt + (2\alpha - 2)|M| \leq (3\alpha - 2)opt$.) Thus the $\frac{4}{3}$ -approximation algorithm of [17] for unweighted 2-ECSS gives a 2-approximation algorithm for MAP. (Although preliminary results and claims have been published on achieving approximation guarantees below $\frac{4}{3}$ for unweighted 2-ECSS, there are no refereed publications to date, see [17].)

Over the last two decades, starting with the work of [15], a few methods have been developed to obtain approximation guarantees below 2 for TAP. The recent methods of [1, 4] rely on so-called bundle constraints defined by paths of zero-cost edges. Unfortunately, these methods (including methods that use the bundle constraints) rely on the fact that the set of zero-cost edges forms a connected graph that spans all the nodes, see [4, 12, 15]. Clearly, this property does not hold for MAP.

1.2 Hardness of approximation of MAP and FAP

MAP is a generalization of the unweighted 2-ECSS problem (consider the special case of MAP with $M = \emptyset$). The latter problem is known to be APX-hard; thus, it has a “hardness of approximation” threshold of $1 + \epsilon$ where $\epsilon > 0$ is a constant, see [7]. Hence, MAP is APX-hard.

Given the lack of progress on approximation algorithms for FAP, one may wonder whether there is a “hardness of approximation” threshold that would explain the lack of progress. Unfortunately, the results and techniques from the area of “hardness of approximation” are far from the known approximation guarantees for many problems in network design. For example, even for the notorious Asymmetric TSP (ATSP), the best “hardness of approximation” lower bound known is around $\frac{75}{74} \approx 1.014$, see [11].

1.3 Our method for MAP

We first present a reduction of any given instance of MAP to a collection of well-structured MAP instances such that the approximation guarantee is preserved, see Sections 2, 3, 4. Then we present a $\frac{7}{4}$ approximation algorithm for a well-structured MAP instance, see Sections 3, 5, 6. Our algorithm starts with a so-called D2 (this is a min-cost 2-edge cover) and then applies ear-augmentation steps. We analyze the cost of the ear-augmentations using an approach similar to the one proposed by Vempala & Vetta for the unweighted 2-ECSS problem [18]. Our presentation is self-contained and formally independent of Vempala & Vetta’s manuscript; also, we address a weighted version of the 2-ECSS problem and our challenge is to improve on the approximation guarantee of 2, whereas Vempala & Vetta’s goal is to achieve an approximation guarantee of $\frac{4}{3}$ for the unweighted 2-ECSS problem.

For the sake of completeness, we have included the proofs of several basic results (e.g., so-called Facts); these should not be viewed as new contributions.

An outline of the paper follows. Section 2 has standard definitions and some preliminary results. Section 3 presents an outline of our algorithm for MAP, and explains what is meant by a well-structured MAP instance. Section 4 presents the pre-processing steps that give an approximation preserving reduction from any instance of MAP to a collection of well-structured MAP instances; some readers may prefer to skip this section (and refer back to the results/details as needed). Sections 5, 6 present the $\frac{7}{4}$ approximation algorithm for well-structured MAP instances, and prove

the approximation guarantee. Section 7 presents examples that give lower bounds on our results on MAP. The first example gives a construction such that $opt \approx \frac{7}{4}\tau$, where τ denotes the minimum cost of a 2-edge cover. The second example gives a construction such that the cost of the solution computed by our algorithm is $\approx \frac{7}{4}opt$.

2 Preliminaries

This section has definitions and preliminary results. Our notation and terms are consistent with [2], and readers are referred to that text for further information.

Let $G = (V, E)$ be a (loop-free) multi-graph with edges of cost either zero or one such that the edges of cost zero form a matching. We take G to be the input graph, and we use n to denote $|V(G)|$. Let M denote the set of edges of cost zero. Throughout, the reader should keep in mind that M is a matching; this fact is used in many of our proofs without explicit reminders. We call an edge of M a *zero-edge* and we call an edge of $E - M$ a *unit-edge*. We call a pair of parallel edges a $\{0, 1\}$ -edge-pair if one of the two edges of the pair has cost zero and the other one has cost one.

We use the standard notion of contraction of an edge, see [16, p.25]: Given a multi-graph H and an edge $e = vw$, the contraction of e results in the multi-graph $H/(vw)$ obtained from H by deleting e and its parallel copies and identifying the nodes v and w . (Thus every edge of H except for vw and its parallel copies is present in $H/(vw)$; we disallow loops in $H/(vw)$.)

For a graph H and a set of nodes $S \subseteq V(H)$, $\delta_H(S)$ denotes the set of edges that have one end node in S and one end node in $V(H) - S$; moreover, $H - S$ denotes $H[V(H) - S]$, the subgraph of H induced by $V(H) - S$. For a graph H and a set of edges $F \subseteq E(H)$, $H - F$ denotes the graph $(V(H), E(H) - F)$. We use relaxed notation for singleton sets, e.g., we use $\delta_H(v)$ instead of $\delta_H(\{v\})$, we use $H - v$ instead of $H - \{v\}$, and we use $H - e$ instead of $H - \{e\}$.

We denote the cost of an edge e of G by $\text{cost}(e)$. For a set of edges $F \subseteq E(G)$, $\text{cost}(F) := \sum_{e \in F} \text{cost}(e)$, and for a subgraph G' of G , $\text{cost}(G') := \sum_{e \in E(G')} \text{cost}(e)$.

For ease of exposition, we often denote an instance G, M by G ; then, we do not have explicit notation for the edge costs of the instance, but the edge costs are given implicitly by $\text{cost} : E(G) \rightarrow \{0, 1\}$, and M is given implicitly by $\{e \in E(G) : \text{cost}(e) = 0\}$. Also, we may not distinguish between a subgraph and its node set; for example, given a subgraph H that contains nodes $v_1, v_2, v_3, v_4, \dots$ we may say that $\{v_1, v_2, v_3\}$ is contained in H .

2.1 2EC, 2NC, bridges and D2

A multi-graph H is called k -edge connected if $|V(H)| \geq 2$ and for every $F \subseteq E(H)$ of size $< k$, $H - F$ is connected. Thus, H is 2-edge connected if it has ≥ 2 nodes and the deletion of any one edge results in a connected graph. A multi-graph H is called k -node connected if $|V(H)| > k$ and for every $S \subseteq V(H)$ of size $< k$, $H - S$ is connected. We use the abbreviations *2EC* for “2-edge connected,” and *2NC* for “2-node connected.”

We assume w.l.o.g. that the input G is 2-edge connected. Moreover, we assume w.l.o.g. that there are ≤ 2 copies of each edge (in any multi-graph that we consider); this is justified since an edge-minimal 2-ECSS cannot have three or more copies of any edge (see Proposition 1 below).

For any instance H , let $opt(H)$ denote the minimum cost of a 2-ECSS of H . When there is no danger of ambiguity, we use opt rather than $opt(H)$.

By a *bridge* we mean a cut edge, i.e., an edge of a connected (sub)graph whose removal results in two connected components, and by a *cut node* we mean a node of a connected (sub)graph whose deletion results in two or more connected components. We call a bridge of cost zero a *zero-bridge* and we call a bridge of cost one a *unit-bridge*.

By a *2ec-block* we mean a maximal connected subgraph with two or more nodes that has no bridges. (Observe that each 2ec-block of a graph H corresponds to a connected component of order ≥ 2 of the graph obtained from H by deleting all bridges.) We call a 2ec-block *pendant* if it is incident to exactly one bridge.

The next result characterizes edges that are not essential for 2-edge connectivity.

Proposition 1. *Let H be a 2EC graph and let $e = vw$ be an edge of H . If $H - e$ has two edge-disjoint v, w paths, then $H - e$ is 2EC.*

By a *2-edge cover* (of G) we mean a set of edges F of G such that each node v is incident to at least two edges of F (i.e., $F \subseteq E(G) : |\delta_F(v)| \geq 2, \forall v \in V(G)$). By $D2(G)$ we mean any minimum-cost 2-edge cover of G (G may have several minimum-cost 2-edge covers, and $D2(G)$ may refer to any one of them); we use $\tau(G)$ to denote the cost of $D2(G)$; when there is no danger of ambiguity, we use $D2$ rather than $D2(G)$, and we use τ rather than $\tau(G)$. Note that $D2$ may have several connected components, and each may have one or more bridges; moreover, if a connected component of $D2$ has a bridge, then it has two or more pendant 2ec-blocks.

The next result follows from Theorem 34.15 in [16, Chapter 34].

Proposition 2. *There is a polynomial-time algorithm for computing $D2$.*

The next result states the key lower bound used by our approximation algorithm.

Fact 3. *Let H be any 2EC graph. Then we have $\text{opt}(H) \geq \tau(H)$.*

By a *bridgeless 2-edge cover* (of G) we mean a 2-edge cover (of G) that has no bridges; note that we have no requirements on the cost of a bridgeless 2-edge cover. We mention that the problem of computing a bridgeless 2-edge cover of minimum cost is NP-hard (there is a reduction from TAP), and there is no approximation algorithm known for the case of nonnegative costs.

2.2 Ear decompositions

An *ear decomposition* of a graph is a partition of the edge set into paths or cycles, P_0, P_1, \dots, P_k , such that P_0 is the trivial path with one node, and each P_i ($1 \leq i \leq k$) is either (1) a path that has both end nodes in $V_{i-1} = V(P_0) \cup V(P_1) \cup \dots \cup V(P_{i-1})$ but has no internal nodes in V_{i-1} , or (2) a cycle that has exactly one node in V_{i-1} . Each of P_1, \dots, P_k is called an *ear*; note that P_0 is not regarded as an ear. We call $P_i, i \in \{1, \dots, k\}$, an *open ear* if it is a path, and we call it a *closed ear* if it is a cycle. An *open ear decomposition* P_0, P_1, \dots, P_k is one such that all the ears P_2, \dots, P_k are open. (The ear P_1 is always closed.)

Proposition 4 (Whitney [19]). *(i) A graph is 2EC iff it has an ear decomposition.*

(ii) A graph is 2NC iff it has an open ear decomposition.

2.3 Redundant 4-cycles

By a *redundant 4-cycle* we mean a cycle C consisting of four nodes and four edges of G such that $V(C) \neq V(G)$, two of the (non-adjacent) edges of C have cost zero, and two of the nonadjacent nodes of C have degree two in G . For example, a 4-cycle $C = u_1, u_2, u_3, u_4, u_1$ with zero-edges u_1u_2, u_3u_4 and unit-edges u_2u_3, u_4u_1 of a 2EC graph G is a redundant 4-cycle provided all edges between $V(C)$ and $V(G) - V(C)$ are incident to either u_1 or u_3 .

Observe that every 2-edge cover of G must contain the edges of every redundant 4-cycle. Also, observe that two different redundant 4-cycles are disjoint (i.e., any node is contained in at most one redundant 4-cycle).

2.4 Bad-pairs and bp-components

For any MAP instance (assumed to be 2EC), we define a *bad-pair* to be a pair of nodes $\{v, w\}$ of the graph such that the edge vw is present and has zero cost, and moreover, the deletion of both nodes v and w results in a disconnected graph. Throughout, unless mentioned otherwise, the term bad-pair refers to a bad-pair of the graph G (in Section 4.4, we discuss bad-pairs of some minors (i.e., subgraphs) of G).

By a *bp-component* we mean one of the connected components resulting from the deletion of a bad-pair from G ; see Figure 1. The set of all bp-components (of all bad-pairs) is a cross-free family, see [16, Chap. 13.4]. Consider two bad-pairs $\{v, w\}$ and $\{y, z\}$. One of the bp-components of $\{v, w\}$ contains $\{y, z\}$; call it C_1 . Then it can be seen that all-but-one of the bp-components of $\{y, z\}$ are contained in C_1 ; the one remaining bp-component of $\{y, z\}$, call it C'_1 , contains $\{v, w\}$ and all of the bp-components of $\{v, w\}$ except C_1 . Thus, the union of the two bp-components C_1 and C'_1 contains $V(G)$.

The following fact is analogous to the fact that every tree on ≥ 2 nodes contains a node v such that all-but-one of the neighbors of v are leaves. (To see this, consider a longest path P of a tree, and take v to be the second node of P .)

Fact 5. *Suppose that G has at least one bad-pair. Then there exists a bad-pair such that all-but-one of its bp-components are free of bad-pairs.*

Remark: Let us sketch an algorithmic proof of Fact 5. For any subgraph G' of G , let $\#bp(G')$ denote the number of bad-pairs (of G) that are contained in $V(G')$. For any bad-pair $\{v, w\}$ and its bp-components C_1, \dots, C_k , let us assume that the indexing is in non-increasing order of $\#bp(C_i)$, thus, we have $\#bp(C_1) \geq \#bp(C_2) \geq \dots \geq \#bp(C_k)$.

We start by computing all the bad-pairs of G . Then we choose any bad-pair $\{v_1, w_1\}$ and compute its bp-components $C_1^{(1)}, \dots, C_{k_1}^{(1)}$. If $\#bp(C_2^{(1)}) = 0$, then we are done; our algorithm outputs $\{v_1, w_1\}$. Otherwise, we pick $C_2^{(1)}$ and choose any bad pair $\{v_2, w_2\}$ contained in $V(C_2^{(1)})$. We compute the bp-components $C_1^{(2)}, \dots, C_{k_2}^{(2)}$ of $\{v_2, w_2\}$ (in G). As stated above, $C_1^{(2)}$ contains $C_1^{(1)}$ and $V(C_1^{(2)})$ contains $\{v_1, w_1\}$; moreover, $\#bp(C_2^{(1)}) > \#bp(C_2^{(2)}) + \dots + \#bp(C_{k_2}^{(2)})$, since the bad-pair $\{v_2, w_2\}$ is contained in $V(C_2^{(1)})$ but $\{v_2, w_2\}$ is disjoint from each of its own bp-components. We iterate these steps (i.e., if $\#bp(C_2^{(2)}) \neq 0$, then we pick any bad-pair $\{v_3, w_3\}$ contained in $V(C_2^{(2)})$, ...), until we find a bad-pair $\{v_\ell, w_\ell\}$ such that $\#bp(C_2^{(\ell)}) = 0$; then we are done. Clearly, this algorithm is correct, and it terminates in $O(n)$ iterations.




Figure 1: Illustration of bad-pairs and bp-components.

2.5 Polynomial-time computations

All of the computations in this paper can be easily implemented in polynomial time, see [16]. We state this explicitly in all relevant results (e.g., Proposition 2, Theorem 6), but we do not elaborate on this elsewhere.

3 Outline of the algorithm

This section has an outline of our algorithm. We start by defining a well structured MAP instance.

Definition 1. *An instance of MAP is called well-structured if it has*

- no $\{0, 1\}$ -edge-pairs,
- no redundant 4-cycles,
- no cut nodes, and
- no bad-pairs.

Section 4 explains how to “decompose” any instance of MAP G into a collection of well-structured MAP instances G_1, \dots, G_k such that a 2-ECSS H of G can be obtained by computing 2-ECSSes H_1, \dots, H_k of G_1, \dots, G_k , and moreover, the approximation guarantee is preserved, i.e., the approximation guarantee on G is \leq the maximum of the approximation guarantees on G_1, \dots, G_k (in other words, $\frac{\text{cost}(H)}{\text{opt}(G)} \leq \max_{i=1, \dots, k} \left\{ \frac{\text{cost}(H_i)}{\text{opt}(G_i)} \right\}$).

Algorithm (outline):

- (0) apply the pre-processing steps (reductions) from Section 4 to obtain a collection of well-structured MAP instances G_1, \dots, G_k ;
for each G_i ($i = 1, \dots, k$), apply steps (1),(2),(3):
- (1) compute $D2(G_i)$ in polynomial time (w.l.o.g. assume $D2(G_i)$ contains all zero-edges of G_i);
- (2) then apply “bridge covering” from Section 5 to $D2(G_i)$ to obtain a bridgeless 2-edge cover \tilde{H}_i of G_i ;
- (3) then apply the “gluing step” from Section 6 to \tilde{H}_i to obtain a 2-ECSS H_i of G_i ;
- (4) finally, output a 2-ECSS H of G from the union of H_1, \dots, H_k by undoing the transformations applied in step (0).

The pre-processing of step (0) consists of four reductions:

- (pp1) handle $\{0, 1\}$ -edge-pairs,
- (pp2) handle redundant 4-cycles,
- (pp3) handle cut nodes, and
- (pp4) handle bad-pairs;

these reductions are discussed in Sections 4.1–4.4. Step (0) applies (pp1) to obtain a collection of MAP instances; after that, there is no further need to apply (pp1), see Fact 10. Then, we

iterate: while the collection of MAP instances has one or more of the latter three “obstructions” (redundant 4-cycles, cut nodes, bad-pairs), we apply (pp2), (pp3), (pp4) in sequence. After $\leq n$ iterations, we have a collection of well-structured MAP instances G_i . Then, we compute an approximately optimal 2-ECSS H_i for each G_i using the algorithm of Theorem 6 (below), and finally, we use the H_i subgraphs to construct an approximately optimal 2-ECSS of G .

Our $\frac{7}{4}$ approximation algorithm for MAP follows from the following key theorem, and the fact that the algorithm runs in polynomial time.

Theorem 6. *Given a well-structured instance of MAP G' , there is a polynomial-time algorithm that obtains a 2-ECSS H' from $D2(G')$ (by adding edges and deleting edges) such that $\text{cost}(H') \leq \frac{7}{4}\tau(G')$.*

We use a credit scheme to prove this theorem; the details are presented in Sections 5 and 6. The algorithm starts with $D2 = D2(G')$ as the current subgraph; we start by assigning $\frac{7}{4}$ initial credits to each unit-edge of $D2$; each such edge keeps one credit to pay for itself and the other $\frac{3}{4}$ credits are taken to be working credits available to the algorithm; the algorithm uses these working credits to pay for the augmenting edges “bought” in steps (2) or (3) (see the outline); also, the algorithm may “sell” unit-edges of the current subgraph (i.e., such an edge is permanently discarded and is not contained in the 2-ECSS output by the algorithm) and this supplies working credits to the algorithm (see Sections 5, 6).

In an ear-augmentation step, we may add either a single ear or a double ear (i.e., a pair of ears); see Section 5 (double ears may be added in Case 3, page 27) and Section 6 (double ears may be added in Case 2, page 31). Although this is not directly relevant, we mention that double ear augmentations are essential in matching theory, see [14, Ch.5.4]. As discussed above, in some of the ear-augmentation steps, we may (permanently) delete some edges from the current subgraph; see Section 5 (edges are deleted in double ear augmentations in Case 3, page 27) and Section 6 (edges are deleted in both Cases 1, 2).

Remark: The following examples show that when we relax the definition of a well-structured MAP instance, then the inequality in Theorem 6 could fail to hold. See Figure 2 for illustrations.

- (a) $\{0, 1\}$ -edge-pairs (i.e., parallel edges of cost zero and one) are present. Then $\frac{\text{opt}}{\tau} \approx 2$ is possible. Our construction consists of a root 2ec-block B_0 , say a 6-cycle of cost 6, and $\ell \gg 1$ copies of the following gadget that are attached to B_0 . The gadget consists of a pair of nodes v, w and two incident edges: a copy of edge vw of cost zero, and a copy of edge wv of cost one. Moreover, we have an edge between v and B_0 of cost one, and an edge between w and B_0 of cost one. Observe that a (feasible) 2-edge cover of this instance consists of B_0 and the two parallel edges (i.e., the two copies of the edge vw) of each copy of the gadget, and it has cost $6 + \ell$. Observe that any 2-ECSS contains the two edges between $\{v, w\}$ and B_0 . Thus, $\text{opt} \geq 2\ell$, whereas $\tau \leq 6 + \ell$.
- (b) Redundant 4-cycles are present. Then $\frac{\text{opt}}{\tau} \approx 2$ is possible. Our construction consists of a root 2ec-block B_0 , say a 6-cycle of cost 6, and $\ell \gg 1$ copies of the following gadget that are attached to B_0 . The gadget consists of a 4-cycle $C = u_1, \dots, u_4, u_1$ that has two zero-edges u_1u_2, u_3u_4 and two unit-edges u_2u_3, u_4u_1 ; moreover, we have an edge between u_1 and B_0 of cost one, and an edge between u_3 and B_0 of cost one. Observe that a (feasible) 2-edge cover of this instance consists of B_0 and the 4-cycle C of each copy of the gadget, and it has cost $6 + 2\ell$. Observe that for any 2-ECSS and for each copy of the gadget, the two edges between C and B_0 as well as the four edges of C are contained in the 2-ECSS. Thus, $\text{opt} \geq 4\ell$, whereas $\tau \leq 6 + 2\ell$.

- (c) Cut nodes are present. Then $\frac{opt}{\tau} \approx 2$ is possible. Our construction consists of ℓ copies of a 3-cycle $C = u_1, u_2, u_3, u_1$ where $u_1 u_2$ is a zero edge and the other two edges have cost one. We “string up” the ℓ copies, i.e., node u_3 of the i th copy is identified with node u_1 of the $(i+1)$ th copy. The optimal solution has all the edges, so $opt = 2\ell$, whereas a (feasible) 2-edge cover consists of a Hamiltonian path together with two more edges incident to the two ends of this path, and it has cost $\ell + 2$; thus $\tau \leq 2 + \ell$.
- (d) Bad-pairs are present. Then $\frac{opt}{\tau} \approx 2$ is possible. An example can be obtained by modifying example (b) above. (Each copy of the gadget is connected to B_0 by two edges that are incident to the nodes u_1 and u_2 instead of the nodes u_1 and u_3 . It can be seen that $opt \geq 4\ell$ whereas $\tau \leq 6 + 2\ell$.)




Figure 2: MAP instances G that are *not* well-structured such that $\frac{opt(G)}{\tau(G)} \approx 2$. Edges of cost zero and one are illustrated by dashed and solid lines, respectively.

4 Pre-processing

This section presents the four reductions used in the pre-processing step of our algorithm, namely, the handling of the $\{0, 1\}$ -edge-pairs, the redundant 4-cycles, the cut nodes, and the bad-pairs.

4.1 Handling $\{0, 1\}$ -edge-pairs

We apply the following pre-processing step to eliminate all $\{0, 1\}$ -edge-pairs. We start with a simple result.

Fact 7. *Let H be an (inclusion-wise) edge-minimal 2EC graph, and let e_1, e_2 be a pair of parallel edges of H . Then $H - \{e_1, e_2\}$ has precisely two connected components, and each of these connected components is 2EC.*

Proof. Let v and w be the end nodes of e_1, e_2 . Observe that $H - \{e_1, e_2\}$ has no v, w path. (Otherwise, H would have three edge-disjoint v, w paths; but then, $H - e_1$ would have two edge-disjoint v, w paths, and by Proposition 1, $H - e_1$ would be 2EC; this would contradict the edge-minimality of H .) It follows that $H - \{e_1, e_2\}$ has precisely two connected components (deleting one edge from a connected graph results in a graph with ≤ 2 connected components). Let C_1, C_2 be the two connected components of $H - \{e_1, e_2\}$; clearly, each of C_1, C_2 contains precisely one of the nodes v, w . Suppose that C_1 is not 2EC; then it has a bridge f . Since the parallel edges e_1, e_2 have exactly one end node in C_1 , f stays a bridge of $C_1 \cup C_2 \cup \{e_1, e_2\}$. This is a contradiction, since $H = C_1 \cup C_2 \cup \{e_1, e_2\}$ is 2EC. \square

Let H be any 2EC graph. We call a $\{0, 1\}$ -edge-pair *essential* if its deletion results in a disconnected graph. When we delete an essential $\{0, 1\}$ -edge-pair then we get two connected components and each one is 2EC, by arguing as in the proof of Fact 7. Hence, when we delete all the essential $\{0, 1\}$ -edge-pairs, then we get a number of connected components C_1, \dots, C_k such that each one is 2EC. Clearly, an approximately optimal 2-ECSS of H can be computed by returning the union of approximately optimal 2-ECSSes of C_1, \dots, C_k together with all the essential $\{0, 1\}$ -edge-pairs of H ; moreover, it can be seen that the approximation guarantee is preserved, that is, the approximation guarantee on H is \leq the maximum of the approximation guarantees on C_1, \dots, C_k .

The next observation allows us to handle the inessential $\{0, 1\}$ -edge-pairs.

Fact 8. *Suppose that H is 2EC and it has no essential $\{0, 1\}$ -edge-pairs. Then there exists a min-cost 2-ECSS of H that does not contain any unit-edge of any $\{0, 1\}$ -edge-pair.*

Proof. Consider a min-cost 2-ECSS H' of H that contains the minimum number of $\{0, 1\}$ -edge-pairs, i.e., among all the optimal subgraphs, we pick one that has the fewest number of parallel edges e, f such that $\text{cost}(e) + \text{cost}(f) = 1$. If H' has no $\{0, 1\}$ -edge-pair, then the fact holds. Otherwise, we argue by contradiction. We pick any $\{0, 1\}$ -edge-pair e, f . Deleting both e, f from H' results in two connected components C_1, C_2 such that each is 2EC, by Fact 7. Now, observe that e, f is not essential for H , hence, $H - \{e, f\}$ has an edge e'' between C_1 and C_2 . We obtain the graph H'' from H' by replacing the unit-edge of e, f by the edge e'' . Clearly, H'' is a 2-ECSS of H of cost $\leq \text{cost}(H')$, and moreover, it has fewer $\{0, 1\}$ -edge-pairs than H' . Thus, we have a contradiction. \square

Now, focus on the input graph G . By the discussion above, we may assume that G has no essential $\{0, 1\}$ -edge-pairs. Then we delete the unit-edge of each $\{0, 1\}$ -edge-pair. By Fact 8, the resulting graph stays 2EC and the optimal value is preserved. Thus, we can eliminate all $\{0, 1\}$ -edge-pairs, while preserving the approximation guarantee.

Proposition 9. *Assume that G has no essential $\{0, 1\}$ -edge-pairs. Let \hat{G} be the multi-graph obtained from G by eliminating all $\{0, 1\}$ -edge-pairs (as discussed above). Then $\text{opt}(G) = \text{opt}(\hat{G})$.*

There is a polynomial-time α -approximation algorithm for MAP on G , if there is such an algorithm for MAP on \hat{G} .

In what follows, we continue to use G to denote the multi-graph obtained by eliminating all $\{0, 1\}$ -edge-pairs (for the sake of notational convenience).

The next fact states that the restriction on the zero-edges of G is preserved when we contract a set of edges E' such that none of the end nodes of the zero-edges in $E - E'$ is incident to an edge of E' (i.e., the end nodes of the zero-edges of G/E' are “original nodes” rather than “contracted nodes.”)

Fact 10. Let $G = (V, E)$ satisfy the restriction on the zero-edges (i.e., G has no $\{0, 1\}$ -edge-pairs, and the zero-edges form a matching). Suppose that we contract a set of edges $E' \subset E$ such that there exists no node that is incident to both an edge in E' and a zero-edge in $E - E'$. Then the restriction on the zero-edges continues to hold for the contracted multi-graph G/E' .

4.2 Handling redundant 4-cycles

We contract all of the redundant 4-cycles in a pre-processing step. Recall that two distinct redundant 4-cycles have no nodes and no edges in common. We first compute all the redundant 4-cycles and then we contract each of these cycles.

Proposition 11. Suppose that G has q redundant 4-cycles. Let \hat{G} be the multi-graph obtained from G by contracting all redundant 4-cycles. Then $\text{opt}(G) = \text{opt}(\hat{G}) + 2q$. There is a polynomial-time α -approximation algorithm for MAP on G , if there is such an algorithm for MAP on \hat{G} .

Remark: Note that the contraction of a redundant 4-cycle may result in “new” cut nodes.

4.3 Handling cut nodes

Let H be any 2EC graph. Then H can be decomposed into blocks H_1, \dots, H_k such that each block is either 2NC or else it consists of two nodes with two parallel edges between the two nodes. (Thus, $E(H)$ is partitioned among $E(H_1), \dots, E(H_k)$ and any two of the blocks H_i, H_j are either disjoint or they have exactly one node in common.) It is well known that an approximately optimal 2-ECSS of H can be computed by taking the union of approximately optimal 2-ECSSes of H_1, \dots, H_k ; moreover, the approximation guarantee is preserved, see [17, Proposition 1.4].

Proposition 12. Suppose that G has cut nodes; let G_1, \dots, G_k be the blocks of G . Then $\text{opt}(G) = \sum_{i=1}^k \text{opt}(G_i)$. There is a polynomial-time α -approximation algorithm for MAP on G , if there is such an algorithm for MAP on $G_i, \forall i \in \{1, \dots, k\}$.

Remark: Note that the blocks G_1, \dots, G_k of a 2EC graph G may contain redundant 4-cycles even if G has no redundant 4-cycle.

4.4 Pre-processing for bad-pairs

This sub-section presents the pre-processing step of our algorithm that handles the bad-pairs. This step partitions the edges of a 2NC instance G among a number of sub-instances G_i such that each sub-instance is 2NC and has no bad-pairs. We ensure the key property that the union of the 2-ECSSes of the sub-instances G_i forms a 2-ECSS of G (see Fact 16).

Throughout this sub-section, unless there is a statement to the contrary, we assume that the instance G has no $\{0, 1\}$ -edge-pairs. This assumption is valid because the pre-processing step (pp1) has been applied already.

4.4.1 Bad-pairs and bp-components

Recall that a *bad-pair* of a 2EC MAP instance is a pair of nodes $\{v, w\}$ such that vw is a zero-edge, and the deletion of both nodes v and w results in a disconnected graph. Throughout, unless mentioned otherwise, the term bad-pair refers to a bad-pair of the graph G .

For an instance \tilde{G} that has no cut nodes and no bad-pairs, we define $\hat{\tau}(\tilde{G})$ to be $2q + \sum_{i=1}^k \tau(\tilde{G}_i)$, where q denotes the number of redundant 4-cycles of \tilde{G} , and $\tilde{G}_1, \dots, \tilde{G}_k$ denotes the collection of well-structured MAP instances obtained from \tilde{G} by repeatedly applying the pre-processing steps (pp2) and (pp3) (observe that (pp2) and (pp3) cannot introduce “new” bad-pairs nor “new” $\{0, 1\}$ -edge-pairs).

Fact 13. *Let \tilde{G} be an instance that has no cut nodes and no bad-pairs (by assumption \tilde{G} has no $\{0, 1\}$ -edge-pairs). We have $\tau(\tilde{G}) \leq \hat{\tau}(\tilde{G}) \leq \text{opt}(\tilde{G})$. Moreover, $\hat{\tau}(\tilde{G})$ can be computed in polynomial time.*

Proof. Let C_1, \dots, C_q denote the redundant 4-cycles of \tilde{G} . Let G' denote the graph obtained from \tilde{G} by contracting C_1, \dots, C_q , and let G'_1, \dots, G'_k denote the blocks of G' . None of G'_1, \dots, G'_k contains a redundant 4-cycle (since no cut node of G' is incident to a zero-edge), hence, each of G'_1, \dots, G'_k is a well-structured MAP instance. Observe that a 2-edge cover of \tilde{G} is given by the union of $\bigcup_{i=1}^k E(\text{D2}(G'_i))$ and $\bigcup_{j=1}^q E(C_j)$, and we have $\text{cost}(E(\text{D2}(G'_i))) = \tau(G'_i)$, $\forall i \in \{1, \dots, k\}$, hence, $\tau(\tilde{G}) \leq (\sum_{i=1}^k \tau(G'_i)) + (2q) = \hat{\tau}(\tilde{G})$. This proves the first inequality.

Let \tilde{G}^{opt} denote an arbitrary min-cost 2-ECSS of \tilde{G} ; clearly, \tilde{G}^{opt} contains each of C_1, \dots, C_q . Let G^* be obtained from \tilde{G}^{opt} by contracting C_1, \dots, C_q . Note that $E(G^*)$ can be partitioned into sets E_1^*, \dots, E_k^* such that $E_i^* \subseteq E(G'_i)$ and E_i^* induces a 2EC subgraph of G'_i , for all $i \in \{1, \dots, k\}$. Hence, $\text{cost}(G^*) = \sum_{i=1}^k \text{cost}(E_i^*) \geq \sum_{i=1}^k \tau(G'_i)$. Therefore, $\text{opt}(\tilde{G}) = \text{cost}(\tilde{G}^{\text{opt}}) = 2q + \text{cost}(G^*) \geq 2q + \sum_{i=1}^k \tau(G'_i) = \hat{\tau}(\tilde{G})$. This proves the second inequality. \square

The pre-processing algorithm and analysis of this subsection rely on using the sharper lower bound of $\hat{\tau}$ on opt .

For a bad-pair $\{v, w\}$ and one of its bp-components C we use $C^{\{v, w\}}$ to denote the subgraph of G induced by $V(C) \cup \{v, w\}$; thus, we have $C^{\{v, w\}} = G[V(C) \cup \{v, w\}]$; moreover, if C has ≥ 2 nodes, then we use C° to denote the multi-graph obtained from $C^{\{v, w\}}$ by contracting the zero-edge vw , whereas if C has only one node then we take C° to be the same as $C^{\{v, w\}}$ (to ensure that C° is 2NC, see Fact 14, we have to ensure that it has ≥ 3 nodes; if C has only one node, then note that $C^{\{v, w\}}/\{vw\}$ has only 2 nodes).

We sketch our plan for handling the bad-pairs in this informal and optional paragraph. Assume that G has one or more bad-pairs. We traverse the “tree of bp-components and bad-pairs”; at each iteration, we pick a bad-pair $\{v_\ell, w_\ell\}$ such that all-but-one of its bp-components are free of bad-pairs, see Fact 5. Let C_1 denote the (unique) bp-component that has one or more bad-pairs (w.l.o.g. assume C_1 exists), and let C_2, \dots, C_k denote the bp-components free of bad-pairs; we call these the *leaf* bp-components. It is easily seen that for a bp-component C_i , the subgraph of any optimal solution induced by $V(C_i) \cup \{v_\ell, w_\ell\}$ has cost $\geq \tau(C_i^\circ)$ (see Fact 15). Now, focus on an optimal solution and let F^* denote its set of edges, i.e., (V, F^*) is a 2-ECSS of G of minimum cost. Since (V, F^*) is 2EC, it can be seen that there exists a $j \in \{1, \dots, k\}$ such that the graph $(C_j^{\{v_\ell, w_\ell\}} - v_\ell w_\ell)$ contains a v_ℓ, w_ℓ path; then, it follows that $F^* \cup \{v_\ell w_\ell\}$ induces a 2-ECSS of $C_j^{\{v_\ell, w_\ell\}}$. In other words, we may assume w.l.o.g. that the zero-edge $v_\ell w_\ell$ (of the bad-pair) is in F^* , and we may “allocate” it to one of the bp-components. Informally speaking, our plan is to return k sub-instances such that one of the sub-instances is of the form $C_j^{\{v_\ell, w_\ell\}}$ while the other sub-instances are of the form C_i° ; this can be viewed as “allocating” the zero-edge $v_\ell w_\ell$ to one carefully chosen bp-component C_j by “mapping” C_j to the sub-instance $C_j^{\{v_\ell, w_\ell\}}$, while all other bp-components C_i ($i \neq j$) are “mapped” to sub-instances C_i° . We “allocate” the zero-edge $v_\ell w_\ell$ as follows: For each $i = 2, \dots, k$, we compute $\hat{\tau}(C_i^{\{v_\ell, w_\ell\}})$ and $\hat{\tau}(C_i^\circ)$. Suppose that there is an index $j \in \{2, \dots, k\}$ such that these two numbers are the same for C_j . Then we “allocate” the zero-edge to C_j ; in case of ties, we pick any j such that $\hat{\tau}(C_j^{\{v_\ell, w_\ell\}}) = \hat{\tau}(C_j^\circ)$. On the other hand,

if the two numbers differ for each $j \in \{2, \dots, k\}$, then we “allocate” the zero-edge to C_1 . (Although this allocation may disagree with the allocation used by the optimal solution, it turns out that we incur no “loss”.) This brings us to the end of the iteration for the bad-pair $\{v_\ell, w_\ell\}$; the algorithm applies the same method to the “remaining graph,” namely, either $C_1^{\{v_\ell, w_\ell\}}$ or C_1° . We mention that C_1 plays a special role in our pre-processing algorithm; this will become clear when we prove its correctness (see Lemma 18 below). See the example in Figure 3.




Figure 3: An example illustrating a “bad allocation” of the zero-edge of a bad-pair. The graph (left subfigure) has one bad-pair $\{a, b\}$ and its bp-components are C_1 (with 4 nodes) and C_2 (with 2 nodes); all edges have cost one, except ab ; note that $opt = 8$. If ab is “allocated” to C_1 , the sum of the costs of the D2 subgraphs of the resulting sub-instances is 9.

Fact 14. *Let G be 2NC (by assumption G has no $\{0, 1\}$ -edge-pairs). Let $\{v, w\}$ be a bad-pair of G , and let C be one of the bp-components of $\{v, w\}$.*

- (i) *Then both $C^{\{v,w\}}$ and C° are 2NC.*
- (ii) *Suppose that G has no redundant 4-cycles and has $\ell \geq 1$ bad-pairs. Then C° has no redundant 4-cycles, and it has $\leq \ell - 1$ bad-pairs, whereas $C^{\{v,w\}}$ has at most one redundant 4-cycle, and it has $\leq \ell - 1$ bad-pairs.*
- (iii) *Suppose that G has no redundant 4-cycles and suppose that C° has no bad-pairs (of its own) and $C^{\{v,w\}}$ has no bad-pairs (of its own). Then $\hat{\tau}(C^\circ) \leq \hat{\tau}(C^{\{v,w\}})$.*

Proof. We prove each of the three parts.

- (i) First, consider $C^{\{v,w\}}$; observe that it has ≥ 3 nodes; since G is 2NC and C is a connected component of $G - \{v, w\}$, for each node $z \in V(C)$, there exist two openly disjoint paths between z and $\{v, w\}$ in $C^{\{v,w\}}$; hence, it can be seen that $C^{\{v,w\}}$ is 2NC.

Now, consider C° , and w.l.o.g., assume that C has ≥ 2 nodes. Then, by definition, C° has ≥ 3 nodes. Let v^* denote the node resulting from the contraction of vw . Arguing as above, for each node $z \in V(C)$, there exist two openly disjoint paths between z and v^* in C° ; moreover, $C = C^\circ - v^*$ has a single connected component; hence, it can be seen that C° is 2NC.

- (ii) Observe that any bad-pair of $C^{\{v,w\}}$ (respectively, C°) is a bad-pair of G . Also, note that $\{v, w\}$ is a bad-pair of G but it is not a bad-pair of $C^{\{v,w\}}$ (since $C = C^{\{v,w\}} - \{v, w\}$ is connected).

Consider C° , and w.l.o.g., assume that C has ≥ 2 nodes. Let v^* denote the node resulting from the contraction of vw . Note that a redundant 4-cycle of C° cannot be incident to v^* (since every edge incident to v^* in C° is a unit-edge). It follows that any redundant 4-cycle of C° is a redundant 4-cycle of G . Clearly, C° has $\leq \ell - 1$ bad-pairs, and it has no redundant 4-cycles; thus, part (ii) holds for C° .

Next, we verify part (ii) for $C^{\{v,w\}}$. Clearly, $C^{\{v,w\}}$ has $\leq \ell - 1$ bad-pairs. Observe that $C^{\{v,w\}}$ has at most one redundant 4-cycle containing the zero-edge vw . Also, observe that any redundant 4-cycle of $C^{\{v,w\}}$ that is disjoint from $\{v, w\}$ is also a redundant 4-cycle of G . It follows that $C^{\{v,w\}}$ has at most one redundant 4-cycle.

(iii) Observe that C° is a well-structured MAP instance, hence, $\widehat{\tau}(C^\circ) = \tau(C^\circ)$.

First, suppose that $C^{\{v,w\}}$ has no redundant 4-cycles. Then, $C^{\{v,w\}}$ is a well-structured MAP instance, hence, $\widehat{\tau}(C^{\{v,w\}}) = \tau(C^{\{v,w\}})$. Moreover, we have $\tau(C^\circ) \leq \tau(C^{\{v,w\}})$; to see this, note that we can start with $D2(C^{\{v,w\}})$ and contract the zero-edge vw to get a 2-edge cover of C° . Therefore, part (iii) holds in this case.

Next, suppose that $C^{\{v,w\}}$ has a redundant 4-cycle. Then, $C^{\{v,w\}}$ has a unique redundant 4-cycle Q that contains the zero-edge vw . Let \widetilde{C} denote the graph obtained from $C^{\{v,w\}}$ by contracting Q . Then, we have $\tau(C^\circ) \leq 2 + \tau(\widetilde{C}) \leq \widehat{\tau}(C^{\{v,w\}})$; the second inequality follows from the definition of $\widehat{\tau}(C^{\{v,w\}})$ (the cost of the union of the D2 subgraphs of the well-structured MAP instances obtained from \widetilde{C} by repeatedly applying (pp2) and (pp3) is $\widehat{\tau}(C^{\{v,w\}}) - 2$); to verify the first inequality note that $E(Q) \cup E(D2(\widetilde{C}))$ is a 2-edge cover of $C^{\{v,w\}}$, and if we contract the zero-edge vw then we get a 2-edge cover of C° . Thus, part (iii) holds in this case. □

See Figure 4 for an illustration.

Remark: Recall that k denotes the number of bp-components of the bad-pair $\{v, w\}$. Readers may focus on the case of $k = 2$ for the rest of this section; our presentation is valid for any $k \geq 2$.

The next fact is essential for our analysis. It states that for any 2-ECSS H of G that contains all the zero-edges, a bad-pair $\{v, w\}$ of G , and any bp-component C_i of $\{v, w\}$, the subgraph of H induced by $V(C_i) \cup \{v, w\}$ is either 2EC or it is connected and has vw as its unique bridge. (To see this, observe that for any node $u \in V(C_i)$, H has two edge-disjoint paths that start at u , end at either v or w , and have all internal nodes in $V(C_i)$.) Therefore, the subgraph obtained from $H[V(C_i) \cup \{v, w\}]$ by contracting the edge vw is 2EC. This implies the key lower bound, $\text{cost}(H[V(C_i) \cup \{v, w\}]) \geq \text{opt}(C_i^\circ)$.

Fact 15. *Let $\{v, w\}$ be a bad-pair, and let C_1, \dots, C_k be all of its bp-components. Let H be a 2-ECSS of G that contains all the zero-edges. Consider any C_i , where $i \in \{1, \dots, k\}$. Then $H[V(C_i) \cup \{v, w\}]$ is connected, and moreover, either it is 2EC or it has exactly one bridge, namely, vw . Therefore, $H[V(C_i^\circ)]$ is 2EC. Hence, $\text{cost}(H[V(C_i) \cup \{v, w\}]) \geq \text{opt}(C_i^\circ)$.*

Fact 16. *Let $\{v, w\}$ be a bad-pair, and let C_1, \dots, C_k be all of its bp-components. Suppose that we pick one of these bp-components C_j and compute a 2-ECSS H_j of $C_j^{\{v,w\}}$. For each of the other bp-components C_i , $i \in \{1, \dots, k\} - \{j\}$, we compute a 2-ECSS H_i of C_i° . Then the union of the edge sets of $H_1, \dots, H_{j-1}, H_j, H_{j+1}, \dots, H_k$ gives a 2-ECSS of G .*

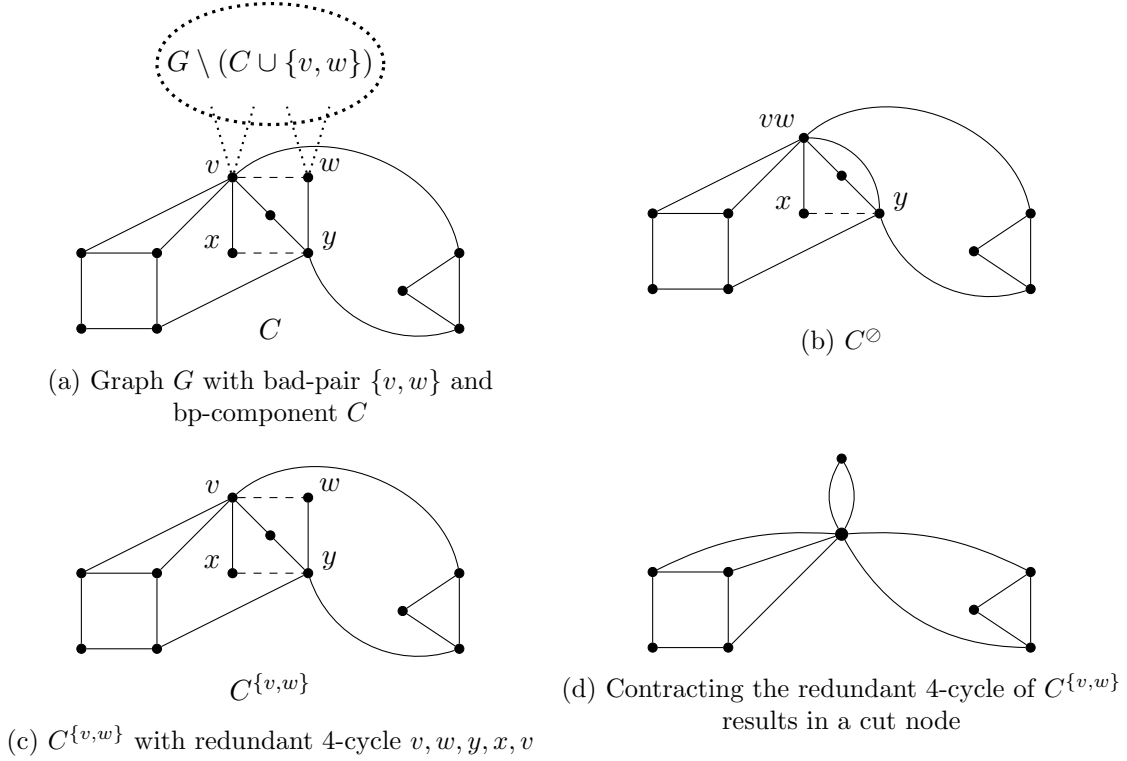


Figure 4: Illustration of a bad-pair $\{v, w\}$ and C^\emptyset and $C^{\{v,w\}}$ for a bp-component C . Dashed lines indicate zero-edges. We have $\hat{\tau}(C^\emptyset) = \tau(C^\emptyset) = 10$; note that $D2(C^\emptyset)$ consists of a 3-cycle and two 4-cycles. Observe that $C^{\{v,w\}}$ has a redundant 4-cycle v, w, y, x, v , and the contraction of this redundant 4-cycle results in a cut node; the resulting graph “decomposes” into three blocks whose D2 subgraphs have costs 2, 4, 5, respectively, hence, $\hat{\tau}(C^{\{v,w\}}) = 13$.

4.4.2 Pre-processing algorithm

Suppose that G is 2NC and it has one or more bad-pairs.

Pre-processing Algorithm (outline):

- (0) pick a bad-pair $\{v, w\}$ that satisfies the condition in Fact 5, and let its bp-components be C_1, C_2, \dots, C_k , where C_2, \dots, C_k are leaf bp-components (free of bad-pairs);
- (1) for each $i = 2, \dots, k$, we compare $\hat{\tau}(C_i^\emptyset)$ and $\hat{\tau}(C_i^{\{v,w\}})$; if the former is strictly smaller than the latter for all $i = 2, \dots, k$, then we return the list of graphs $C_2^\emptyset, \dots, C_k^\emptyset$ and $C_1^{\{v,w\}}$; informally speaking, we allocate the zero-edge vw to C_1 ;
- (2) otherwise, we have at least one $j \in \{2, \dots, k\}$ such that $\hat{\tau}(C_j^\emptyset) = \hat{\tau}(C_j^{\{v,w\}})$; then we return the list of graphs $C_1^\emptyset, \dots, C_{j-1}^\emptyset, C_{j+1}^\emptyset, \dots, C_k^\emptyset$ and $C_j^{\{v,w\}}$; informally speaking, we allocate the zero-edge vw to C_j ;
- (3) let G' denote the “remaining graph,” either $C_1^{\{v,w\}}$ or C_1^\emptyset ; **stop** if G' has no bad-pairs, otherwise, apply the same pre-processing iteration to G' .

For any instance of MAP G' , we use $lb(G')$ to denote

$$\begin{cases} \hat{\tau}(G'), & \text{if } G' \text{ has no cut nodes and no bad-pairs (and no } \{0,1\}\text{-edge-pairs),} \\ \text{opt}(G'), & \text{otherwise.} \end{cases}$$

Fact 17. *Suppose that a 2-ECSS of cost $\leq \alpha \cdot \tau(G_i)$ can be computed in polynomial time for any well-structured MAP instance G_i . Then for any MAP instance G' that has no cut nodes and no bad-pairs (and no $\{0,1\}$ -edge-pairs), a 2-ECSS of cost $\leq \alpha \cdot \hat{\tau}(G')$ can be computed in polynomial time.*

Proof. By the pre-processing, we “decompose” G' into a collection of well-structured MAP instances G'_1, \dots, G'_q ; then, for each G'_i , we compute a 2-ECSS H'_i of cost $\leq \alpha \cdot \tau(G'_i)$; finally, we undo the transformations applied by the pre-processing to obtain a 2-ECSS H' of G' from the collection H'_1, \dots, H'_q . Our previous arguments (recall the proof of Fact 13) imply that $\text{cost}(H') \leq \alpha \cdot \hat{\tau}(G')$. \square

Lemma 18. *Let G be a MAP instance that has no cut-nodes and has $\rho \geq 1$ bad-pairs (by assumption G has no $\{0,1\}$ -edge-pairs). Let G_1, \dots, G_k denote the collection of graphs obtained by applying one iteration of the above pre-processing algorithm for a bad-pair $\{v, w\}$ of G (thus, $\{v, w\}$ satisfies the condition in Fact 5, and each G_i is of the form $C_i^{\{v,w\}}$ or C_i°). Then we have*

$$lb(G) \geq lb(G_1) + \sum_{i=2}^k \hat{\tau}(G_i);$$

moreover, G_1 has $\leq \rho - 1$ bad-pairs, and G_2, \dots, G_k have no bad-pairs.

Proof. Let C_1, \dots, C_k denote all of the bp-components of $\{v, w\}$; note that C_2, \dots, C_k are leaf bp-components (free of bad-pairs). For each $i = 1, \dots, k$, let V_i denote the set of nodes $V(C_i) \cup \{v, w\}$; thus, V_i denotes the node-set of $C_i^{\{v,w\}}$.

Let G^* denote an optimal solution, i.e., a 2-ECSS of minimum cost, and w.l.o.g. assume that G^* contains all zero-edges of G . We have $\text{opt}(G) = \text{cost}(G^*) = \sum_{i=1}^k \text{cost}(G^*[V_i])$.

By Fact 15, we have

$$\text{if } G_i \text{ is of the form } C_i^\circ, \text{ then } \text{cost}(G^*[V_i]) \geq \text{opt}(C_i^\circ) \geq \hat{\tau}(G_i), \quad \forall i \in \{2, \dots, k\}. \quad (*)$$

We complete the proof by examining a few cases.

Case 1: the zero-edge vw is allocated to C_1 : First, suppose that $G^*[V_1]$ is 2EC. Then, we have $\text{cost}(G^*[V_1]) \geq \text{opt}(G_1) \geq lb(G_1)$; combining this with (*) we have

$$\text{opt}(G) = \sum_{i=1}^k \text{cost}(G^*[V_i]) \geq \sum_{i=1}^k \text{opt}(G_i) \geq \sum_{i=1}^k lb(G_i).$$

Now, suppose that $G^*[V_1]$ is not 2EC; then, by Fact 15, it is connected and has only one bridge, namely, vw . Moreover, G has an edge $\hat{e} \neq vw$ whose end nodes are in two different connected components of $G^*[V_1] - vw$. (To see this, note that $G[V_1]$ is 2NC, hence, $(G[V_1] - vw)$ has a v, w path, and one of the edges of this path satisfies the requirement on \hat{e} .) Thus, adding \hat{e} to $G^*[V_1]$ results in a 2EC graph of cost $1 + \text{cost}(G^*[V_1]) \geq \text{opt}(G_1) \geq lb(G_1)$. Also, observe that G^* has two edge-disjoint v, w paths, hence, one of the subgraphs $G^*[V_\ell] - vw$ has a v, w path, where $\ell \in \{2, \dots, k\}$. Hence,

$$\text{cost}(G^*[V_\ell]) \geq \text{opt}(C_\ell^{\{v,w\}}) \geq \hat{\tau}(C_\ell^{\{v,w\}}) \geq 1 + \hat{\tau}(C_\ell^\circ) = 1 + \hat{\tau}(G_\ell)$$

(we used the fact that $\widehat{\tau}(C_i^{\{v,w\}}) > \widehat{\tau}(C_i^\circ)$, $\forall i \in \{2, \dots, k\}$). By the above inequalities and (*), we have

$$\text{opt}(G) = \sum_{i=1}^k \text{cost}(G^*[V_i]) \geq \underbrace{lb(G_1) - 1 + (\widehat{\tau}(G_\ell) + 1)}_{i=1} + \sum_{i \in \{2, \dots, \ell-1, \ell+1, \dots, k\}} \widehat{\tau}(G_i) = lb(G_1) + \sum_{i=2}^k \widehat{\tau}(G_i).$$

Case 2: the zero-edge vw is allocated to C_ℓ , $\ell \in \{2, \dots, k\}$: Thus, we have $\widehat{\tau}(C_\ell^\circ) = \widehat{\tau}(C_\ell^{\{v,w\}})$ (otherwise, vw cannot be allocated to C_ℓ). Then, by Fact 15, we have

$$\text{cost}(G^*[V_\ell]) \geq \text{opt}(C_\ell^\circ) \geq \widehat{\tau}(C_\ell^\circ) = \widehat{\tau}(C_\ell^{\{v,w\}}).$$

This, together with (*) and the inequalities $\text{cost}(G^*[V_1]) \geq \text{opt}(G_1) \geq lb(G_1)$, implies that $\text{opt}(G) \geq lb(G_1) + \sum_{i=2}^k \widehat{\tau}(G_i)$.

This completes the proof of the lemma. □

Theorem 19. *Suppose that a 2-ECSS of cost $\leq \alpha \cdot \tau(G_i)$ can be computed in polynomial time for any well-structured MAP instance G_i . Then for any MAP instance G , a 2-ECSS of cost $\leq \alpha \cdot lb(G)$ can be computed in polynomial time via the pre-processing (consisting of (pp1) followed by iterations of (pp2), (pp3), (pp4)).*

Proof. We use induction on the number of bad-pairs in the MAP instance G .

Suppose that G has no bad-pairs. Then, by the pre-processing, we “decompose” G into a collection of well-structured MAP instances G_1, \dots, G_q ; then, for each G_i , we compute a 2-ECSS H_i of cost $\leq \alpha \cdot \tau(G_i)$; finally, we undo the transformations applied by the pre-processing to obtain a 2-ECSS H of G from the collection H_1, \dots, H_q . It follows from the results in Sections 4.1–4.3 that $\text{cost}(H) \leq \alpha \cdot lb(G)$.

Now, suppose that the theorem holds for any MAP instance with $\leq \ell$ bad-pairs.

Suppose that G has $\ell + 1$ bad-pairs. For notational convenience in the induction step, let us assume that G has no $\{0, 1\}$ -edge-pairs and no cut nodes. (When we have an arbitrary instance of MAP, G , then we apply (pp1) and (pp3) to “decompose” the instance into a collection of MAP instances that each have no $\{0, 1\}$ -edge-pairs and no cut nodes; then, we apply the following proof to each of the “restricted” MAP instances; finally, we undo the transformations applied by (pp1) and (pp3), and we apply the results in Sections 4.1 and 4.3 to obtain a 2-ECSS of G of cost $\leq \alpha \cdot lb(G)$.)

Then, by one iteration of (pp4) and Lemma 18, we “decompose” G into a collection of MAP instances G_1, \dots, G_k such that

- G_1 has no $\{0, 1\}$ -edge-pairs, no cut nodes, and $\leq \ell$ bad-pairs, and G_2, \dots, G_k have no $\{0, 1\}$ -edge-pairs, no cut nodes, and no bad-pairs,
- $lb(G) \geq lb(G_1) + \sum_{i=2}^k \widehat{\tau}(G_i)$,
- given a 2-ECSS H_i of G_i for each $i \in \{1, \dots, k\}$, the union of H_1, \dots, H_k forms a 2-ECSS H of G , by Fact 16.

By the induction hypothesis and by Fact 17 (which is essential for this analysis), we may assume that $\text{cost}(H_1) \leq \alpha \cdot lb(G_1)$ and $\text{cost}(H_i) \leq \alpha \cdot \widehat{\tau}(G_i)$, $\forall i \in \{2, \dots, k\}$. Then, H is a 2-ECSS of G , and $\text{cost}(H) \leq \alpha \cdot lb(G)$ (since $lb(G_1) + \sum_{i=2}^k \widehat{\tau}(G_i) \leq lb(G)$). □

5 Bridge covering

In this section and in Section 6, we assume that the input is a well-structured MAP instance.

We start by illustrating our method for bridge covering on a small example. After computing D2, recall that we give 1.75 initial credits to each unit-edge of D2, thereby giving each of these edges 0.75 working credits (we “retain” one credit to buy the unit-edge for our solution). Consequently, each 2ec-block of D2 gets ≥ 1.5 working credits (this is explained below). We want to buy more edges to add to D2 such that all bridges are “covered”, and we pay for the newly added edges via the working credits.

Observe that each 2ec-block of D2 has ≥ 1.5 working credits, because it has ≥ 2 unit-edges; to see this, suppose that a 2ec-block B has b nodes; if $b = 2$, then B has two parallel unit-edges; otherwise, B has $\geq b$ edges (since B is 2EC) and has $\leq \lfloor b/2 \rfloor$ zero-edges (since the zero-edges form a matching), so B has $\geq \lceil b/2 \rceil$ unit-edges, and we have $b \geq 3$.

In what follows, we use the term credits to mean the working credits of the algorithm; this excludes the unit credit retained by every unit-edge of the current solution subgraph; for example, a 6-cycle of D2 that contains four unit-edges has 3 credits (and it has 7 initial credits).

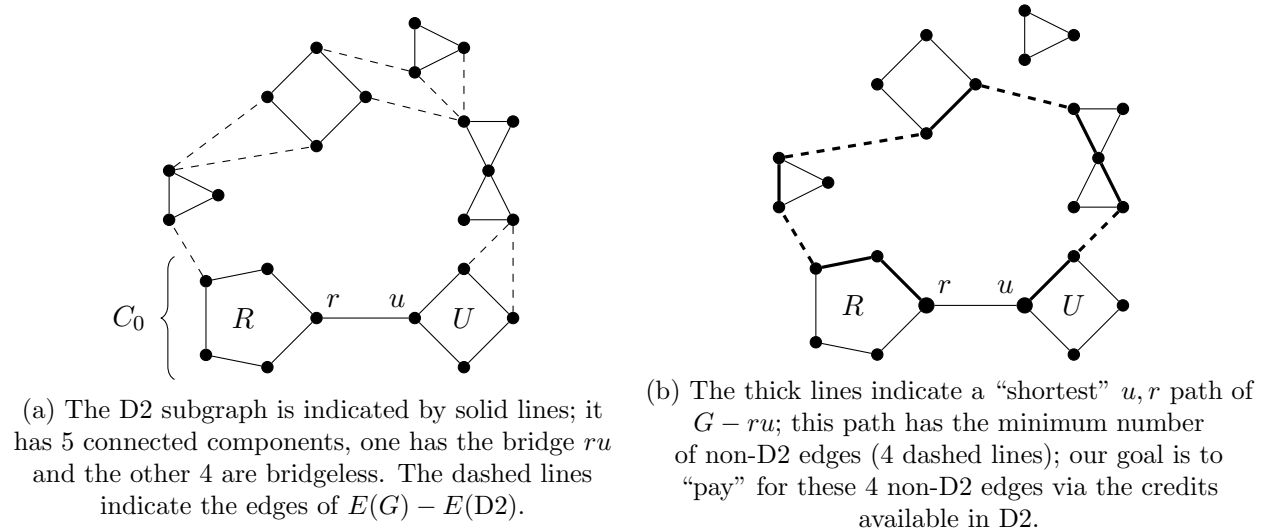


Figure 5: Illustration of bridge covering on a simple example.

Consider an example such that D2 has a connected component C_0 that has one bridge and two 2-ec blocks R and U ; let ru denote the unique bridge where r is in R and u is in U . (It can be seen that ru is a zero-bridge, but we will not use this fact.) Since G is assumed to be 2NC, $G - ru$ is connected, hence, it contains a u, r path; let us pick a u, r path Y of $G - ru$ that has only its prefix and suffix in common with C_0 and that has the minimum number of non-D2 edges. Our plan is to augment D2 by adding the edge set $E(Y) - E(D2)$, thus “covering” the bridge ru . We may view this as an “ear-augmentation step” that adds the ear** Y . We have to pay for the non-D2 edges of Y by using the credits available in D2. Let us traverse Y from u to r , and each time we see a non-D2 edge of Y , then we will pay for this edge. For the sake of illustration, consider the example in Figure 5; note that the u, r path Y (in Figure 5(b), on the right) has $\ell = 4$ non-D2 edges (indicated by dashed lines). When we traverse Y starting from u , then observe that each of the first $(\ell - 1)$ of these edges has its last node in a distinct 2ec-block of D2 (moreover, none of

**Throughout, we abuse the term *ear*; although Y is not an ear, one may view the minimal subpath of Y from U to R , call it Y' , as an ear of G w.r.t. C_0 , i.e., Y' is a path of G that has both end nodes in C_0 and has no internal node in C_0 .

these $(\ell - 1)$ 2ec-blocks is in C_0). We pay for these $(\ell - 1)$ edges by borrowing one credit from the credit of each of these $(\ell - 1)$ 2ec-blocks. We need one more credit to pay for the last non-D2 edge of Y . We get this credit from the prefix of Y between its start and its first non-D2 edge; in particular, we borrow one credit from the credit of U . Thus, we can pay for all the non-D2 edges of Y . Observe that by adding the edge set $E(Y) - E(D2)$ to D2, we have merged several 2ec-blocks (including R and U) into a new 2ec-block. We give the new 2ec-block (that contains R and U) the credit of R as well as any unused credit of the other 2ec-blocks incident to Y .

The general case of bridge covering is more complicated.

Remark: For the sake of exposition, we may impose a direction on a path, cycle, or ear (e.g., we traverse Y from u to r in the discussion above). Nevertheless, the input G is an undirected graph, so (formally speaking) there is no direction associated with paths or cycles of G .

5.1 Post-processing D2

Immediately after computing D2, we apply a post-processing step that replaces some unit-edges of D2 by other unit-edges to obtain another D2 that we denote by $\widehat{D2}$ that satisfies the following key property:

Every pendant 2ec-block B of $\widehat{D2}$ that is incident to a zero-bridge has $\text{cost}(B) \geq 3$, and hence, has ≥ 2.25 credits (see part (2) of the credit invariant below).

In other words, if a 2ec-block of $\widehat{D2}$ has ≤ 2 unit-edges, then either the 2ec-block is not pendant (i.e., it is incident to no bridges or ≥ 2 bridges) or it is pendant and is incident to a unit-bridge.

For any subgraph G' of G , let $F_0(G')$ denote the set of zero-bridges (of G') that are incident to pendant 2ec-blocks (of G') of cost ≤ 2 (the notation F_0 is used only in this subsection); moreover, let $\#\text{comp}(G')$ denote the number of connected components of G' . Thus, the goal of the post-processing step is to compute $\widehat{D2}$ such that $F_0(\widehat{D2})$ is empty.

The post-processing step is straightforward. W.l.o.g. assume that the initial D2 contains all the zero-edges; we start with this D2 and iterate the following step. Let $D2^{old}$ denote the D2 at the start of the iteration. If $F_0(D2^{old})$ is empty then we are done, we take $D2^{old}$ to be $\widehat{D2}$. Otherwise, we pick any zero-bridge v_0v_1 in $F_0(D2^{old})$, and we take B to be a pendant 2ec-block with $\text{cost}(B) = 2$ that is incident to v_0v_1 ; note that B exists by the definition of F_0 . Now, observe that all of the edges of B incident to v_0 have cost one, hence, it can be seen that B has either 2 or 3 nodes; then, since G is 2NC and $D2^{old}$ contains all the zero-edges, G has a unit-edge e between $V(B) - \{v_0\}$ and $V - V(B)$; moreover, B has a unit-edge f that is incident to e and to v_0 . We replace f by e . The resulting subgraph is a 2-edge cover of the same cost; we denote it by $D2^{new}$.

We claim that $F_0(D2^{new}) \subseteq F_0(D2^{old})$, and moreover,

$$|F_0(D2^{new})| + \#\text{comp}(D2^{new}) < |F_0(D2^{old})| + \#\text{comp}(D2^{old}).$$

To see this, first suppose that the end nodes of e are in two different connected components of $D2^{old}$; then it can be seen that $F_0(D2^{new}) \subseteq F_0(D2^{old})$, and hence, the claimed inequality holds; otherwise, (the end nodes of e are in the same connected component of $D2^{old}$) v_0v_1 is not a bridge of $D2^{new}$, therefore, $F_0(D2^{new})$ is a proper subset of $F_0(D2^{old})$, and hence, the claimed inequality holds.

Thus, after $O(n)$ iterations, we find $\widehat{D2}$ that satisfies the key property. See Figure 6 for illustrations.

Proposition 20. *There is a polynomial-time algorithm that finds a 2-edge cover $\widehat{D2}$ (of cost $\tau(G)$) such that $F_0(\widehat{D2})$ is empty; thus, $\widehat{D2}$ satisfies the key property.*

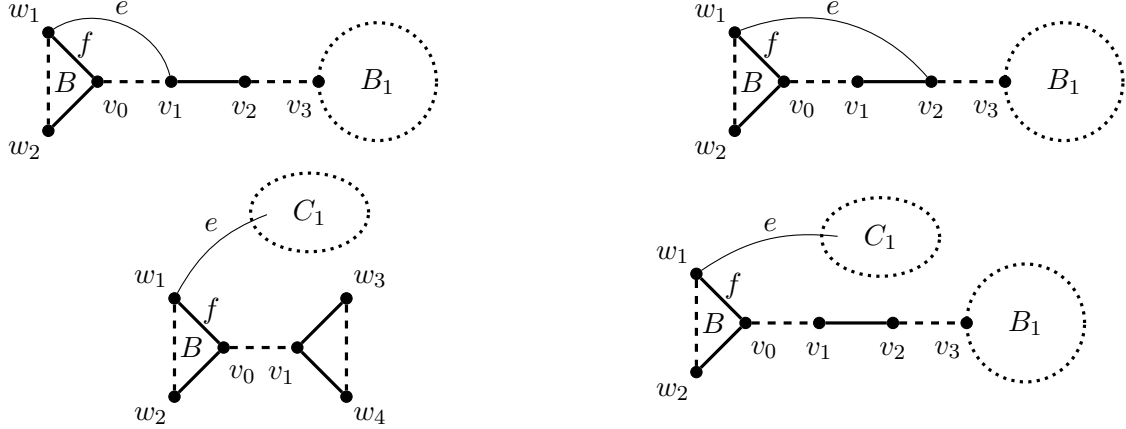


Figure 6: Illustrations of some cases that could arise during post-processing. Edges of D_2 are indicated by thick lines. Edges of cost zero and one are indicated by dashed and solid lines, respectively.

5.2 Credit invariant and charging lemma

Let $H = (V, F)$ denote the current graph of picked edges; thus, at the start, H is the same as $\widehat{D_2}$, and we may assume that H has one or more bridges (otherwise, bridge covering is trivial). By an *original* 2ec-block B of H we mean a 2ec-block (of the current H) such that B is also a 2ec-block of $\widehat{D_2}$ (i.e., the set of edges incident to $V(B)$ is the same in both $\widehat{D_2}$ and the current H). By a *new* 2ec-block of H we mean a 2ec-block (of the current H) that is not an original 2ec-block. Similarly, we define an original/new connected component of H .

For any subgraph H' of H , we use $\text{credits}(H')$ to denote the sum of the working credits of the unit-edges of H' . At the start of bridge covering, $\text{credits}(H)$ is equal to the sum of the initial credits of the unit-edges of $\widehat{D_2}$ minus the number of unit-edges of $\widehat{D_2}$.

We call a node v of H a *white* node if v belongs to a 2ec-block of H , otherwise, we call v a *black* node.

Fact 21. *Suppose that H has one or more black nodes. Let v be a black node of H . Then all edges of H incident to v are bridges of H , and v is incident to ≥ 2 bridges of H . Every maximal path of H that starts with v contains a white node.*

Remark: Let H' denote the graph obtained from H by contracting each 2ec-block; thus, each 2ec-block of H maps to a “contracted” white node of H' , each black node of H maps to a black node of H' , and each bridge of H maps to a bridge of H' . Clearly, H' is a forest; it may have isolated nodes (these correspond to 2EC connected components of H). Clearly, H' has ≥ 2 edges incident to each black node.

By a *b-path* of H we mean a path consisting of bridges that starts and ends with arbitrary nodes (white or black) such that all internal nodes are black nodes. We say that two 2ec-blocks of a connected component of H are *b-adjacent* if there exists a b-path whose terminal nodes are in these two 2ec-blocks.

Initially, our algorithm picks a connected component C_0 of $H = \widehat{D_2}$ that has one or more bridges. If C_0 has any pendant 2ec-block that is incident to a zero-bridge, then the algorithm picks such a 2ec-block and designates it as the *root 2ec-block* R (in this section, R always denotes the root 2ec-block), otherwise, the algorithm picks any pendant 2ec-block of C_0 and designates it as the root

2ec-block R . Also, the algorithm picks the unique bridge of C_0 incident to R ; we denote this bridge by ru , where r is in R . If ru is a zero-bridge, then since C_0 is an original connected component, R has ≥ 2.25 credits (by the key property of $\widehat{D2}$); otherwise, R has ≥ 1.5 credits. Immediately after designating R , we ensure that R has ≥ 2 credits. If R is short of credits (i.e., it has only 1.5 credits), then we borrow 0.5 credits from a 2ec-block of C_0 that is b-adjacent to R (in this case, observe that every pendant 2ec-block of C_0 is incident to a unit-bridge); see credit invariant (3) below.

Each iteration of bridge covering maintains the following invariant, i.e., if the invariant holds at the start of an iteration, then it holds at the end of that iteration. (At the end of this section, we argue that this invariant is preserved at the end of every iteration, see Proposition 25. Note that the invariant could be temporarily violated within an iteration, due to various updates.)

Credit invariant:

- (1) *Each unit-edge of H has a unit of retained credit (recall that retained credits are distinct from working credits, and the term “credit” means working credit). Each original 2EC connected component has ≥ 1.5 credits, and each new 2EC connected component has ≥ 2 credits. Moreover, each unit-bridge of H has 0.75 credits.*
- (2) *Within each original connected component of H , each pendant 2ec-block that is incident to a zero-bridge has ≥ 2.25 credits.*
- (3) *Suppose that the root 2ec-block R is well defined. Then either R is incident to a zero-bridge and has ≥ 2.25 credits, or R has ≥ 2 credits. Moreover, each 2ec-block that is b-adjacent to R has ≥ 1 credits, and every other 2ec-block of H has ≥ 1.5 credits.*

In an arbitrary iteration of the algorithm, either R is a pendant 2ec-block of an original connected component C_0 of H , or R has been designated as the root 2ec-block by the previous iteration and there exists a bridge (of H) incident to R . In the former case, (as discussed above) the algorithm chooses ru to be the unique bridge of C_0 incident to R ; in the latter case, if there exists a unit-bridge incident to R , then the algorithm chooses ru to be such a bridge, otherwise, the algorithm chooses ru to be any zero-bridge incident to R .

When we remove the edge ru from C_0 then we get two connected components; let us denote them by C_0^r (it contains r but not u) and C_0^u (it contains u but not r). Recall that $G - ru$ has a path between C_0^u and C_0^r . Let P be such a path of $G - ru$ that starts at some node \hat{a} of C_0^u , ends at some node \hat{z} of C_0^r , has no internal nodes in C_0 , and (subject to the above) minimizes $|E(P) - E(H)|$. Note that there could be many choices for the nodes \hat{a} and \hat{z} ; although the choice of these nodes is important for our analysis (see Lemma 23 and its proof), the next lemma (Lemma 22) and its proof apply for all valid choices of these two nodes. See Figure 7.




Figure 7: Illustration of plan for “bridge covering.” Large circles indicate 2ec-blocks. Dashed lines indicate P . Q is the path of C_0 between r and \hat{a} , where \hat{a} is the unique node of P in C_0^u .

Let H^{new} denote $H \cup (E(P) - E(H))$ and let B^{new} denote the 2ec-block of H^{new} that contains \hat{a} , \hat{z} , and R . We designate B^{new} as the root 2ec-block of H^{new} , provided B^{new} is incident to a bridge of H^{new} .

Each iteration may be viewed as an ear-augmentation step that adds either one open ear to C_0 or two open ears to C_0 , e.g., the path P may be viewed as an open ear of G w.r.t. C_0 . On the other hand, our charging scheme (for paying for the edges of $E(P) - E(H)$) views each iteration as adding an ear w.r.t. R , that is, we take the ear to be the union of three paths, namely, an r, \hat{a} path of C_0 that contains ru , the path P , and a path of C_0 between \hat{z} and R (we mention that our charging scheme also uses the credits available in the first of these three paths). To avoid confusion, we call these the R -ears, and unless mentioned otherwise, an “ear” means an ear w.r.t. C_0 .

Let us outline our credit scheme (for bridge covering) and its interaction with an arbitrary iteration of the algorithm. Let $\hat{a} \in V(C_0^u)$, $\hat{z} \in V(C_0^r)$, and P be as described above. Let C_1, \dots, C_k denote the connected components of H that contain (at least) one internal node of P (thus, $C_0 \neq C_i$, $\forall i = 1, \dots, k$). Lemma 22 and its proof, see below, explain how the credits of C_1, \dots, C_k can be redistributed such that each of C_1, \dots, C_k releases one unit of credit such that the credit invariant holds again at the end of the iteration. Thus, we can buy all-but-one of the unit-edges of $E(P) - E(H)$ using the credits released by C_1, \dots, C_k . In order to establish the credit invariant at the end of the iteration, we have to find another unit of credit (to pay for one unit-edge of $E(P) - E(H)$), and, in addition, we have to ensure that credit invariant (3) is maintained (e.g., we may have to find another 0.25 units of credit). We defer this issue to Section 5.3; in fact, this is the most intricate part of bridge covering.

Lemma 22. *Let H, C_0, C_0^u, C_0^r be as stated above, and suppose that H satisfies the credit invariant. Let P be an open ear w.r.t. C_0 with end nodes \hat{a}, \hat{z} , where \hat{a} is in C_0^u and \hat{z} is in C_0^r , such that each connected component of H that contains an internal node of P is incident to exactly two edges of $E(P) - E(H)$ (i.e., $E(P) - E(H)$ forms a path in the auxiliary graph obtained from G by contracting all connected components of H other than C_0). Let C_1, \dots, C_k denote the connected components of H that contain (at least) one internal node of P (thus, $C_0 \neq C_i$, $\forall i = 1, \dots, k$). Then the credits of C_1, \dots, C_k can be redistributed such that each of C_1, \dots, C_k releases one unit of credit such that the credit invariant holds again at the end of the iteration.*

More precisely, if C_i ($i \in \{1, \dots, k\}$) has a 2ec-block that contains a node of P , then we take one unit from the credit of one such 2ec-block B . (At the end of the iteration, B is “merged into” B^{new} , hence, the credit invariant holds again.) Otherwise, we take 0.5 units from the credits of two distinct 2ec-blocks B_1, B_2 of C_i that are b-adjacent to P . (At the end of the iteration, B_1 and B_2 are b-adjacent to B^{new} , hence, the credit invariant holds again.)

Proof. Our goal is to show that we can pay via the credits of C_1, \dots, C_k for all-but-one of the edges of $E(P) - E(H)$, while ensuring that the credit invariant holds again for H^{new} . When we traverse P from \hat{a} to \hat{z} , observe that each of the edges of $E(P) - E(H)$, except the last such edge, has its last end node in a distinct connected component $C_i \neq C_0$ of H ; we use the credits available in that connected component to pay the cost of the edge. The rest of the proof shows how we can obtain one unit from the credit of the relevant connected component while ensuring that the credit invariant continues to hold for H^{new} . (Proposition 25 below shows that the credit invariant is preserved at the end of each iteration.)

Consider any original connected component $C \neq C_0$ of H that contains one of the internal nodes of P (thus C is one of C_1, \dots, C_k). See Figure 8. By our choice of P , there is a unique edge of P that “enters” C and there is a unique edge of P that “exits” C , i.e., there is a unique edge f with one end node in C and the other end node in the subpath of P between \hat{a} and C , and




Figure 8: Illustration of the notation in Lemma 22. A connected component $C \neq C_0$ of H is shown (C does not contain R). The large circles indicate 2ec-blocks of C . The path P is indicated by dashed lines; P contains 3 edges of C . The arrows indicate maximal b-paths ending at 2ec-blocks.

similarly, there is a unique edge e with one end node in C and the other end node in the subpath of P between C and \hat{z} .

Let s_0 denote the end node of f in C , and let t_0 denote the end node of e in C . Possibly, $s_0 = t_0$. Let $P(s_0, t_0)$ denote the s_0, t_0 sub-path of P . Clearly, $P(s_0, t_0)$ is contained in H .

First, suppose that $P(s_0, t_0)$ contains a white node; there is a 2ec-block of C , call it B , that contains this white node; then we take one unit from the credit of B and use that to pay for the edge f . (Recall that B^{new} is designated as the root 2ec-block of H^{new} , and note that B^{new} contains both s_0 and t_0 ; moreover, the 2ec-block B (of H) is also contained in B^{new} ; hence, at the end of the iteration, it can be seen that the credit invariant is maintained in H^{new} although we borrowed one credit from B ; see the proof of Proposition 25.)

Otherwise, both s_0 and t_0 are black nodes. Then we resort to a more complex scheme. Let $C(s_0, s_1)$ be any maximal b-path of $C - E(P(s_0, t_0))$ that starts with s_0 and ends with a white node s_1 , and let $B(s_1)$ denote the 2ec-block that contains the white node s_1 ; note that $B(s_1)$ has no nodes in common with B^{new} . Similarly, let $B(t_1)$ denote a 2ec-block that contains the terminal white node t_1 (where, $t_1 \neq s_1$) of a maximal b-path of $C - E(P(s_0, t_0))$ that starts with t_0 ; it can be seen that $B(t_1)$ has no nodes in common with B^{new} .^{††} We take 0.5 credits from each of $B(s_1)$ and $B(t_1)$ and use that to pay for the edge f . (In H^{new} , observe that both s_0 and t_0 are contained in the root 2ec-block B^{new} ; moreover, $B(s_1)$ and $B(t_1)$ are 2ec-blocks, and moreover, both these 2ec-blocks are b-adjacent to B^{new} ; hence, at the end of the iteration, the credit invariant is maintained in H^{new} although we borrowed 0.5 credits from each of $B(s_1)$ and $B(t_1)$; see the proof of Proposition 25.)

Finally, consider any new connected component $C \neq C_0$ of H that contains one of the internal nodes of P . Our algorithm ensures that every new connected component of H , except for C_0 , is 2EC, hence, C is 2EC. We take one unit from the credit of C and use that to pay for the unique edge of P that “enters” C . (At the end of the iteration, the credit invariant is maintained in H^{new} because C is contained in B^{new} , the designated root 2ec-block of H^{new} ; see the proof of Proposition 25.) \square

5.3 Algorithm and analysis for bridge covering

We present the algorithm and analysis of bridge covering based on Lemma 22.

Recall that H satisfies the credit invariant initially, and that Lemma 22 allows us to pay for all-but-one of the unit-edges added by a single ear-augmentation. Our goal is to charge the remaining

^{††}If $s_0 = t_0$, then note that s_0 is incident to ≥ 2 bridges of $C = C - E(P(s_0, t_0))$, hence, we can ensure that $t_1 \neq s_1$.

cost (of one) to a prefix of the R -ear that we denote by Q ; Q is the maximal path of C_0 contained in the R -ear and starting with the edge ru . Let a_0 denote the other end node of Q (thus, when we traverse the edges (and nodes) of the R -ear starting with the edge ru , then a_0 is the first node incident to an edge of the R -ear in $E(G) - E(H)$). Since our goal is to collect as much credit as possible from Q we choose the R -ear such that either (i) Q has a white node w ($w \neq r$) or (ii) Q has no white nodes (other than r), Q has the maximum cost possible, and subject to this, Q has the maximum number of bridges possible. In other words, we choose (the 3-tuple) P, a_0, z_0 such that P is a path of $G - ru$ with one end node a_0 in C_0^u and the other end node z_0 in C_0^r , none of the internal nodes of P is in C_0 , the associated prefix Q satisfies condition (i) or (ii) (stated above), and (subject to all the above requirements) P has the minimum number of edges from $E(G) - E(H)$. We can easily compute (the 3-tuple) P, a_0, z_0 in polynomial time via standard methods from graph algorithms; this is discussed in the next lemma.

Lemma 23. *P, a_0, z_0 satisfying the requirements stated above can be computed in polynomial time.*

Proof. For each node $v \in C_0^u$, we define $\gamma(v)$ as follows: $\gamma(v) = \infty$ if every path of C_0 between v and r contains a white node w , $w \neq r$; otherwise, $\gamma(v) = |E(C_0(v, r))| + n \cdot \text{cost}(C_0(v, r))$, where $C_0(v, r)$ denotes the unique b-path of C_0 between v and r . (Informally speaking, $\gamma(v)$ assigns a “rank” to each node v of C_0^u ; if every v, r path of C_0 contains two or more white nodes, then the rank is ∞ , otherwise, the rank is determined by the unique b-path of C_0 between v and r , and we rank according to the 2-tuple consisting of the cost and the number of bridges of the relevant path.)

Then, we construct the following weighted directed graph: the directed graph has two oppositely oriented edges for each edge of $G - V(C_0)$, it has an edge oriented out of $V(C_0^r)$ for each edge of $G - ru$ in the cut $\delta_G(V(C_0^r))$, and it has an edge oriented into $V(C_0^u)$ for each edge of $G - ru$ in the cut $\delta_G(V(C_0^u))$ (there are no oriented edges corresponding to other edges of G). We assign weights of zero to the oriented edges associated with the edges of H , and weights of one to the other oriented edges (associated with the edges of $E(G) - E(H)$). Then, we apply a reachability computation, taking all the nodes in C_0^r to be the sources. We claim that a node $v \in C_0^u$ is reachable from C_0^r (in the directed graph) iff $G - ru$ has a path between C_0^r and v such that no internal node of the path is in C_0 .

Thus, we can find a_0 and z_0 that satisfy the requirements: we choose a_0 to be a node $v \in C_0^u$ that is reachable from C_0^r (in the directed graph) and that has the maximum $\gamma()$ value, and then we choose z_0 to be a node in C_0^r such that the directed graph has a path from this node to a_0 . Then, we take P to be a shortest z_0, a_0 path in the (weighted) directed graph. \square

An outline of the bridge covering step follows.

Bridge covering (outline):

- (0) compute D2, then post-process D2 to obtain $\widehat{D2}$, and let $H = \widehat{D2}$;
- (1) pick any (original) connected component C_0 of H that has a bridge; if possible, pick a pendant 2ec-block of C_0 that is incident to a zero-bridge and designate it as the root 2ec-block R ; otherwise, (every pendant 2ec-block of C_0 is incident to a unit-bridge) pick any pendant 2ec-block of C_0 and designate it as the root 2ec-block R ;
- (2) **repeat**
 - if possible, pick a unit-bridge of C_0 incident to R , otherwise, pick any zero-bridge of C_0 incident to R , and denote the picked bridge by ru ;
 - apply one ear-augmentation step (add one or two ears, see Cases 1–3 below) to cover a sub-path of bridges starting with ru , and let B^{new} denote the resulting 2ec-block that contains R and ru ;
 - let $R := B^{new}$;

until $R = B^{new}$ has no incident bridges of H ;
- (3) **stop** if H has no bridges, otherwise, **go to** (1).

Recall from Section 5.2 that we have to establish the credit invariant at the end of the iteration, hence, we have to find another unit of credit (to pay for one unit-edge of $E(P) - E(H)$), and, in addition, we have to ensure that credit invariant (3) is maintained (e.g., we may have to find another 0.25 units of credit).

Lemma 24. *Let H, C_0, R, P, Q, B^{new} be as stated above, and suppose that H satisfies the credit invariant. Except for one case, $\text{credits}(Q) + \text{credits}(R)$ suffices to give one unit of credit for $E(P) - E(H)$ and to give sufficient credit to B^{new} such that the credit invariant holds at the end of the iteration.*

In the exceptional case (described in Case 3, see below), the algorithm adds another ear P_ with associated prefix Q_* such that $\text{credits}(Q_* - Q) \geq 0.75$ and, moreover, no connected component of $H - C_0$ is incident to both P and P_* , and “sells” a unit-bridge of Q (i.e., a unit-bridge of $\widehat{D2}$ is permanently discarded and is not contained in the 2-ECSS output by the algorithm, thereby releasing 1.75 credits). Thus, the algorithm obtains ≥ 2.5 credits, and this suffices to give one unit of credit for each of $E(P) - E(H)$ and $E(P_*) - E(H)$, and to give sufficient credit to B^{new} such that the credit invariant holds at the end of the iteration.*

The rest of this subsection presents a proof of this lemma. Before presenting the details, we give an informal overview of the credit scheme in the following two tables. The goal is to find one unit of credit for each ear added by the current iteration, and ensure that B^{new} has sufficient credits at the end of the iteration, assuming that the credit invariant holds at the start of the iteration. If our iteration adds one ear, then clearly it suffices to find 1.25 credits (but there are cases where the goal can be achieved with fewer credits). We have two cases: either Q has a white node w other than r , or else Q has no white node other than r .

Properties of Q	Credits available to algorithm excluding $\text{credits}(R)$
Q has a white node w other than r	≥ 1.0 (from 2ec-block containing w)
subcase: Q contains a unit-bridge	≥ 1.75 (1.0 from above & 0.75 from unit-bridge)
subcase: Q contains no unit-bridges	≥ 1.0 (1.0 from above; note: R has ≥ 2.25 credits)

Properties of Q	Credits available to algorithm excluding credits(R)
Q contains no white nodes other than r , Q has ≥ 2 bridges, $\text{cost}(Q) \geq 1.0$, node sequence of Q : $v_0 = r, v_1 = u, v_2, \dots, v_k = a_0$	
case: $\text{cost}(Q) \geq 2$	≥ 1.5 (this suffices)
case: $\text{cost}(Q) = 1.0$	(see 4 subcases below)
subcase: Q has 2 bridges, thus $a_0 = v_2$, and $\text{cost}(v_0v_1) = 1, \text{cost}(v_1v_2) = 0$	not possible
subcase: Q has 2 bridges, thus $a_0 = v_2$, and $\text{cost}(v_0v_1) = 0, \text{cost}(v_1v_2) = 1$, and H has a unit-bridge v_2v_3	$= 0.75$ (this suffices)
subcase: Q has 2 bridges, thus $a_0 = v_2$, and $\text{cost}(v_0v_1) = 0, \text{cost}(v_1v_2) = 1$, and the only other bridge of H incident to v_2 is a zero-bridge v_2v_3	
update H : add another ear P_\star , permanently discard unit-bridge v_1v_2	≥ 2.5 (this suffices)
subcase: Q has 3 bridges, thus $a_0 = v_3$, and $\text{cost}(v_0v_1) = 0, \text{cost}(v_1v_2) = 1, \text{cost}(v_2v_3) = 0$	$= 0.75$ (this suffices)

Suppose that the prefix Q of the R -ear contains a white node w other than r (thus, $w \neq r$). Let $B \neq R$ denote the 2ec-block that contains w ; clearly, B has at least one credit (see credit invariant (3)). Observe that B^{new} contains B and the credits of both B and Q are available. If Q has at least one unit-bridge, then the sum of the credits of B and Q is ≥ 1.75 , and this suffices. Now, suppose that Q has no unit-bridges. Thus, Q has only one bridge, namely, the zero-bridge ru . Then, by our choice of ru and the credit invariant, R already has ≥ 2.25 credits, hence, we need only one credit for the ear-augmentation step (since B^{new} gets all the credits of R). Thus, $\text{credits}(B) = 1$ suffices when Q has no bridge other than ru .

Now, assume that the prefix Q has no white node other than r ; in particular, the nodes u and a_0 are black. It is easily seen that Q has ≥ 2 bridges. (Since G is 2NC, $G - u$ has a path between the connected component of $C_0 - u$ that contains R and each of the other connected components of $C_0 - u$; this implies that there exists an R -ear such that its associated prefix Q' has ≥ 2 bridges and has $\text{cost}(Q') \geq 1$; then, by our choice of Q , we have $\text{cost}(Q) \geq \text{cost}(Q') \geq 1$ and Q has ≥ 2 bridges.)

Suppose that $\text{cost}(Q) \geq 2$, thus, Q has ≥ 2 unit-bridges. Then, $\text{credits}(Q) \geq 1.5$. This suffices to pay for the current ear-augmentation step and to re-establish the credit invariant.

Now, we may assume that $\text{cost}(Q) = 1$. Let us denote the node sequence of Q by $v_0, v_1, v_2, \dots, v_k$, where $v_0 = r, v_1 = u$, and $v_k = a_0$. Clearly, we have three possibilities:

Case 1: Q consists of 2 bridges and $\text{cost}(v_0v_1) = 1, \text{cost}(v_1v_2) = 0$: We argue that this possibility cannot occur. Observe that $v_2 = a_0$. Observe that $G - \{v_1, v_2\}$ is connected; otherwise, $\{v_1, v_2\}$ would be a bad-pair. Consider the connected components of $C_0 - \{v_1, v_2\}$; let S denote the set of nodes of the connected component that contains R , and let T denote $V(C_0) - \{v_1, v_2\} - S$. Since $G - \{v_1, v_2\}$ is connected, it has path P_\star between a node $z_\star \in S$ and a node $a_\star \in T$ such that P_\star has no internal nodes in C_0 . Observe that every path of C_0 between r and a_\star has $\text{cost} \geq 2$ (because such a path contains the unit-edge ru as well

as another edge of the cut $\delta_{C_0}(\{v_1, v_2\})$, and all edges of this cut have cost one); thus, the prefix Q_\star associated with P_\star has cost ≥ 2 . This contradicts our choice of P, a_0, z_0 (because the prefix Q has cost one). See Figure 9.




Figure 9: Illustration of Case 1 (bridge covering).

Case 2: Q consists of 3 bridges and $\text{cost}(v_0v_1) = 0$, $\text{cost}(v_1v_2) = 1$, $\text{cost}(v_2v_3) = 0$: Then, we have 0.75 credits available in Q . We argue that this suffices to pay for the current ear-augmentation step and to re-establish the credit invariant. Observe that $v_3 = a_0$. By our choice of ru and the credit invariant, R has ≥ 2.25 credits. Thus, we can pay one unit for one unit-edge of $E(P) - E(H)$, and we have 2 credits (from R) left for B^{new} . Moreover, v_3 is a black node, and it can be seen that one of the unit-bridges of C_0 incident to v_3 becomes a bridge incident to B^{new} ; in other words, at the next iteration, when we designate B^{new} as the root 2ec-block, then only 2 credits are required for B^{new} . See Figure 10.




Figure 10: Illustration of Case 2 (bridge covering).

Case 3: Q consists of 2 bridges and $\text{cost}(v_0v_1) = 0$, $\text{cost}(v_1v_2) = 1$: Observe that $v_2 = a_0$. Note that v_2 is a black node, and it is incident to a bridge other than v_2v_1 . There are two subcases, namely, either v_2 is incident to two unit-bridges or not, and we choose v_3 appropriately. In the former subcase, we take v_2v_3 to be a unit-bridge, and in the latter subcase we have to take v_2v_3 to be the unique zero-bridge incident to v_2 .

We handle the first subcase (with $\text{cost}(v_2v_3) = 1$) similarly to Case 2 above. By our choice of ru and the credit invariant, R has ≥ 2.25 credits. Thus, we can pay one unit for one unit-edge of $E(P) - E(H)$, and we have 2 credits (from R) left for B^{new} . Moreover, v_2 is a black node, and it can be seen that the unit-bridge v_2v_3 becomes a bridge incident to B^{new} ; in other words, at the next iteration, when we designate B^{new} as the root 2ec-block, then only 2 credits are required for B^{new} .

In the second subcase (with $\text{cost}(v_2v_3) = 0$), observe that H has precisely two bridges incident to v_2 , because v_2 is incident to only one unit-bridge. Our plan is to add a second ear and then observe that the edge v_1v_2 becomes redundant after the addition of the two ears, hence, we can permanently discard this edge from H thereby gaining 1.75 credits. (Note that when we permanently discard a unit-edge of $\widehat{D2}$ from our solution subgraph H , then all of the retained credits and the working credits of that edge become available.) Moreover, we get another 0.75 (or more) credits from the addition of the two ears. Thus we get ≥ 2.5 credits,

and this suffices to pay for the addition of two ears and to re-establish the credit invariant. See Figure 11.




Figure 11: Illustration of Case 3 (bridge covering).

Observe that $G - \{v_2, v_3\}$ is connected; otherwise, $\{v_2, v_3\}$ would be a bad-pair. Consider the connected components of $C_0 - \{v_2, v_3\}$; let S denote the set of nodes of the connected component that contains R , and let T denote $V(C_0) - \{v_2, v_3\} - S$. Since $G - \{v_2, v_3\}$ is connected, it has a path P_* between a node $z_* \in S$ and a node $a_* \in T$ such that P_* has no internal nodes in C_0 . Moreover, w.l.o.g. we assume that P_* has the minimum number of edges from $E(G) - E(H)$. It can be seen that P_*, a_*, z_* satisfy the following:

- (i) Either $z_* = v_1$ or there is a bridge v_1w_1 (of C_0) such that $v_1w_1 \neq v_1v_0$, $v_1w_1 \neq v_1v_2$, and z_* is in the connected component of $C_0 - v_1w_1$ that contains w_1 . This follows from a contradiction argument; the only other possibility is that z_* is in C_0^r (the connected component of $C_0 - ur$ that contains r); but then we would define the prefix \tilde{Q}_* associated with P_* to be a path of C_0 between r and a_* ; note that \tilde{Q}_* would contain v_3 (since C_0 has only two bridges incident to v_2), hence, \tilde{Q}_* would have ≥ 3 bridges and we would have $\text{cost}(\tilde{Q}_*) \geq 1$; this would contradict our choice of P, a_0, z_0 .
- (ii) Now, we define the prefix Q_* associated with P_* to be a path of C_0 between v_1 and a_* . Note that Q_* contains v_3 (since C_0 has only two bridges incident to v_2). We have two cases: either v_3 is a black node or it is a white node; in the first case, $Q_* - Q$ contains a unit-bridge incident to v_3 so we have $\text{credits}(Q_* - Q) \geq 0.75$, whereas in the second case, $Q_* - Q$ contains the white node v_3 so we can obtain ≥ 1 credit from the 2ec-block (of C_0) that contains v_3 .
- (iii) There is no connected component of $H - C_0$ that is incident to both P and P_* . Otherwise, suppose that some connected component \hat{C} of H is incident to both P and P_* . Then there is a path in $(P \cup P_* \cup \hat{C})$ that starts at z_0 and ends at a_* , and whose prefix $Q_* \cup \{v_0v_1\}$ in C_0 either has a white node or has $\text{cost} \geq 2$, contradicting our choice of P, a_0, z_0 .

Hence, we can apply Lemma 22 separately to each of P_* and P and thus pay for all-but-one of the unit-edges added by each of the two ears.

- (iv) After adding the two ears, the edge v_1v_2 can be permanently discarded from H while preserving 2-edge connectivity, by Proposition 1. To see this, observe that $(C_0 - v_1v_2) \cup P$ contains a v_1, v_2 path, and also $(C_0 - v_1v_2) \cup P_*$ contains a v_1, v_2 path, and moreover, these two paths have no internal nodes in common (i.e., $(C_0 - v_1v_2) \cup P \cup P_*$ contains a cycle incident to v_1 and v_2).
- (v) Consider the credits available from Q and $Q_* - Q$ for the double ear-augmentation. ‘‘Selling’’ the unit-edge v_1v_2 gives 1.75 credits (since v_1v_2 is a unit-bridge of $\widehat{D2}$, the sum

of its retained credit and working credit is 1.75, and since v_1v_2 is permanently discarded from H , all of this credit is released). Moreover, we have ≥ 0.75 credits available in $Q_\star - Q$. Thus, ≥ 2.5 credits are available.

- (vi) By modifying the arguments in the proof of Lemma 23, we can compute (a 3-tuple) $P_\star, a_\star, z_\star$ that has the required properties in polynomial time. Hence, a double ear-augmentation can be implemented in polynomial time.

Summarizing, when v_2v_3 is a zero-edge, then we add the two ears P, P_\star and permanently discard the edge v_1v_2 from H ; we have sufficient credits to pay for the addition of the two ears and to re-establish the credit invariant.

Proposition 25. (i) *The credit invariant holds for $\widehat{D2}$.* (ii) *The credit invariant holds at the end of every iteration of bridge covering (i.e., every iteration of bridge covering preserves the credit invariant).*

Proof. It is easily seen that (i) holds.

Now, focus on any iteration. We use the notation of this section; in particular, we use P to denote the first ear added in an iteration; moreover, we use \widehat{P} to denote the corresponding R -ear (see page 22); also, let R^{new} denote the root 2ec-block of H^{new} (assume it exists).

It is easy to verify that parts (1) and (2) of the credit invariant are preserved by every iteration. Whenever we add a unit-edge e to H (in an ear-augmentation) then we ensure that e has a unit of retained credit, by Lemmas 22, 24. Whenever we take away credits (either working credit or retained credit) from a unit-bridge e of C_0 (e.g., when e is in the prefix Q of \widehat{P}), then we retain sufficient credits for e (if e stays in H^{new} then it keeps ≥ 1 retained credit and ≥ 0 working credit, otherwise, it keeps ≥ 0 retained credit and ≥ 0 working credit); moreover, if we take away credit from a bridge e of C_0 , then either e is contained in the 2ec-block R^{new} of H^{new} or else e is permanently discarded from H (thus, at most one iteration can take away credit from a unit-bridge of C_0).

Consider part (3) of the credit invariant. First, let us consider the credits of R and R^{new} . Suppose that R has α credits at the start of an iteration. Then, aside from three exceptions, at the start of the next iteration, R^{new} has $\geq 0.5 + \alpha$ credits (see the tables placed after the statement of Lemma 24). The first exception occurs when Q has a white node w other than r and Q consists of one zero-bridge; in this case, $\text{credits}(R^{new}) = \text{credits}(R) \geq 2.25$. The other two exceptions occur in cases 2 and 3(first subcase); in these two cases, R has $\alpha \geq 2.25$ credits (since ru is a zero-bridge) while R^{new} has $\geq \alpha - 0.25 \geq 2$ credits and R^{new} is guaranteed to be incident to a unit-bridge of H^{new} . Hence, the credit invariant pertaining to R is preserved.

Consider any single ear-augmentation. Consider the credits of the other (non-root) 2ec-blocks of H and H^{new} . Any 2ec-block of H that contains a node of \widehat{P} is “merged” into R^{new} , hence, the credit invariant is not relevant for such 2ec-blocks of H (we argued above that R^{new} satisfies the credit invariant). Consider a 2ec-block B of H such that there is a b-path (of H) between B and a black node of \widehat{P} . Then, in H^{new} , R^{new} is b-adjacent to B , hence, B is required to have ≥ 1 credit (by the credit invariant). Although we may remove credits from such 2ec-blocks (e.g., see the proof of Lemma 22, and note that we take away 0.5 credits from each of $B(s_1)$ and $B(t_1)$), we ensure that each such 2ec-block has at least one credit at the next iteration. Lastly, consider any 2ec-block B of H that is (node) disjoint from $R \cup \widehat{P}$ and is not b-adjacent to any black node of \widehat{P} . Clearly, B has the same credit in both H and H^{new} . It follows that credit invariant (3) is preserved in every single ear-augmentation.

Now, consider a double ear-augmentation; thus, we have the second subcase of case 3 that adds the two ears P, P_* . Let \widetilde{P}_* denote the union of the R -ear corresponding to P and the R -ear corresponding to P_* . The arguments in the previous paragraph can be re-applied with one change: we replace \widehat{P} by \widetilde{P}_* . It follows that credit invariant (3) is preserved in every double ear-augmentation. \square

This concludes the discussion of bridge covering.

Proposition 26. *At the termination of the bridge covering step, H is a bridgeless 2-edge cover and the credit invariant holds (thus, every original 2ec-block of H has ≥ 1.5 credits and every new 2ec-block of H has ≥ 2 credits). The bridge covering step can be implemented in polynomial time.*

6 The gluing step

In this section and in Section 5, we assume that the input is a well-structured MAP instance.

In this section, we focus on the last step of the algorithm, namely, the gluing step. Our goal here is to show that the credits in H suffice to update H to a 2-ECSS of G by adding some edges and deleting some edges (i.e., the difference between the number of edges added and the number of edges deleted in the gluing step is \leq the credit of H at the start of the gluing step). The following result summarizes this section:

Proposition 27. *At the termination of the bridge-covering step, let H denote the bridgeless 2-edge cover computed by the algorithm and suppose that the credit invariant holds; let γ denote $\text{credits}(H)$. Then the gluing step augments H to a 2-ECSS H' of G (by adding edges and deleting edges) such that $\text{cost}(H') \leq \text{cost}(H) + \gamma$. The gluing step can be implemented in polynomial time.*

It is convenient to define the following multi-graph: let \widehat{G} be the multi-graph obtained from G by contracting each 2ec-block B_i of H into a single node that we will denote by B_i . Observe that \widehat{G} is 2EC. Note that the algorithm “operates” on G and never refers to \widehat{G} ; but, for our discussions and analysis, it is convenient to refer to \widehat{G} .

At the start of the gluing step, recall that each original 2ec-block of H has ≥ 1.5 credits and each new 2ec-block of H has ≥ 2 credits. We pick any 2ec-block R_0 of H and designate it as the root R ; then we apply iterations; each iteration adds to H the edges of an ear whose start node and end node are in R ; some iterations may add a second ear. After each iteration, we update the notation so that R denotes the 2ec-block of the current subgraph H that contains R_0 . Our plan is to keep adding ears to H until all of the nodes are in R . Thus, we terminate when H is 2EC, and on termination H is a 2-ECSS of G . In the following discussion, we assume that R has zero credits, i.e., we ignore the credits available in R while paying for the edges added in ear-augmentation steps.

Consider a cycle \widehat{C} of \widehat{G} incident to R ; such a cycle exists since \widehat{G} is 2EC. Let $|\widehat{C}|$ denote the number of edges of \widehat{C} ; clearly, $|\widehat{C}| = |V(\widehat{C})|$. Observe that each of the non-root nodes of \widehat{C} has ≥ 1.5 credits, hence, we have $\geq 1.5(|\widehat{C}| - 1)$ credits available from (the nodes of) \widehat{C} and this is $\geq |\widehat{C}|$ whenever $|\widehat{C}| \geq 3$. Thus, we have enough credit to buy all the edges of \widehat{C} whenever $|\widehat{C}| \geq 3$. Moreover, if $|\widehat{C}| = 2$ and the non-root node of \widehat{C} has ≥ 2 credits, then we have enough credit to buy all the edges of \widehat{C} . Thus, we have insufficient credit only when $|\widehat{C}| = 2$ and the non-root node of \widehat{C} has exactly 1.5 credits.

In what follows, we assume that all the cycles of \widehat{G} incident to R have insufficient credit. Let \widehat{C} be a cycle of \widehat{G} that has insufficient credit, where $\widehat{C} = R, B, R$ and B denotes the non-root node of \widehat{C} . Clearly, B is an original 2ec-block of H (otherwise, it would have 2 credits rather than 1.5 credits). Let b denote the number of nodes^{‡‡} of B . We have $b \in \{2, 3, 4\}$ because B has

^{‡‡}From this sentence till the end of Section 6, “node” means a node of G (and not a node of \widehat{G}).

$\geq \lceil b/2 \rceil$ unit-edges (B has $\geq b$ edges by 2EC and has $\leq \lfloor b/2 \rfloor$ zero-edges), hence, $\text{credits}(B) < 2$ implies $b < 5$. Moreover, for $b \in \{3, 4\}$, note that B cannot have a cut node (otherwise, B would have $\geq 1 + \lceil b/2 \rceil$ unit-edges), and hence, (since B is 2NC) B must contain a spanning cycle. For $b \in \{3, 4\}$, let $Q(B)$ denote any spanning cycle of B .

The two edges of \hat{C} correspond to two edges of G between B and R ; let e denote one of these edges, and let v_e denote the end node of e in B . Since G is 2NC, $G - v_e$ has a path between $(B - v_e)$ and R . Each such path has all its internal nodes in $V(R) \cup (V(B) - \{v_e\})$ (by our assumption on cycles of \hat{G} incident to R), hence, there exists an edge f of G between $(B - v_e)$ and R ; let u_f denote the end node of f in $B - v_e$. See Figure 12.




Figure 12: Illustration of R , B , and v_e, u_f .

Now, we have two cases, depending on whether v_e and u_f are adjacent in B or not.

Case 1: e, f can be chosen such that B has an edge between v_e and u_f : In this case, we claim that e, f can be chosen such that B has a unit-edge between v_e and u_f . By way of contradiction, suppose that $v_e u_f$ is a zero-edge. Then we have $b \in \{3, 4\}$ (otherwise, if $b = 2$, then B would consist of two parallel unit-edges), and moreover, $G - \{v_e, u_f\}$ is connected (since $\{v_e, u_f\}$ is not a bad-pair). Hence, G has an edge rv_0 such that r is in R and v_0 is in $B - \{v_e, u_f\}$ (we also used our assumption on cycles of \hat{G} incident to R). Moreover, the spanning cycle $Q(B)$ has an edge between v_0 and $\{v_e, u_f\}$. W.l.o.g. suppose that B has the edge $v_0 u_f$; this is a unit-edge (since $v_e u_f$ is a zero edge). Then by replacing the pair of edges e, f by rv_0, f , we have two edges between R and B such that their end nodes in B are distinct and there exists a unit-edge of B between these two end nodes. Our claim follows.

Now, observe that the graph $H \cup \{e, f\} - \{v_e u_f\}$ has two edge-disjoint v_e, u_f paths. We buy the edges e, f and permanently discard the unit-edge $v_e u_f$; that is, we add the two edges e, f to H and remove the edge $v_e u_f$ from H . (In the gluing step, when we permanently discard a unit-edge from our solution subgraph H , then one unit of retained credit of that edge become available.) In the resulting graph H^{new} , the connected component containing R_0 (as well as R and B) is 2EC, by Proposition 1. This step results in a surplus of 0.5 credits (we get 1.5 credits from B , one credit from selling $v_e u_f$, and we pay two credits for the edges e, f).

Case 2: for any choice of e, f there is no edge between v_e and u_f in B : Then, clearly $b = |V(B)| = 4$, and B has a spanning cycle $Q(B)$. Let $Q = v_1, v_2, v_3, v_4, v_1$ denote $Q(B)$, where w.l.o.g. $v_e = v_1$ and $u_f = v_3$. Since B has 1.5 credits, two of the (non-adjacent) edges of Q must be zero-edges. There must be one or more edges of G incident to v_2 or v_4 , otherwise, Q would be a redundant 4-cycle of G .

Suppose that G has the edge $v_2 v_4$. See Figure 13(a). We buy the three edges $e, f, v_2 v_4$ and we permanently discard the two unit-edges of Q . In the resulting graph H^{new} , the connected component containing R_0 (as well as R and B) is 2EC, by Proposition 1. This step results in a surplus of 0.5 credits (we get 1.5 credits from B , two credits from selling the two unit edges of Q , and we pay three credits for the edges $e, f, v_2 v_4$).

Lastly, suppose that v_2 and v_4 are nonadjacent in G . Then G has an edge between another 2ec-block B' of H (where $B' \neq B$ and $B' \neq R$) and one of v_2 or v_4 , say v_2 ; let us denote this

edge by \tilde{e} . Since G is 2NC, there is a path in $G - \{v_2\}$ between B' and $B - \{v_2\}$. Let \tilde{P} denote such a path that has the fewest edges of $E(G) - E(H)$. In \hat{G} , observe that $\tilde{P} \cup \{\tilde{e}\}$ corresponds to a cycle $\hat{C}_{B'}$ that is incident to B and B' . Moreover, in \hat{G} , note that $\hat{C}_{B'}$ cannot be incident to R (by our assumption on cycles of \hat{G} incident to R). See Figure 13(b). Let e_Q denote the unit-edge of Q that has its end nodes among v_1, v_2, v_3 . It can be seen that $H \cup \{e, f, \tilde{e}\} \cup E(\tilde{P}) - \{e_Q\}$ has two edge-disjoint paths between the end nodes of e_Q . We buy e, f , we permanently discard e_Q , and moreover, we buy the edges of $(E(\tilde{P}) - E(H)) \cup \{\tilde{e}\}$. In the resulting graph H^{new} , the connected component containing R_0 (as well as R and B) is 2EC, by Proposition 1. (For example, suppose that $e_Q = v_1v_2$ and the end node of \tilde{P} in B is v_3 ; then, we add the ear formed by e, v_1v_4, v_4v_3, f , and after that, we add the edges in $E(G) - E(H)$ of the closed ear formed by $v_3v_2, \tilde{e}, \tilde{P}$.) This step results in a surplus of credits; note that the sum of the credits of the 2ec-blocks (excluding B) incident to \tilde{P} minus the size of $(E(\tilde{P}) - E(H)) \cup \{\tilde{e}\}$ is at least -0.5 .




Figure 13: Illustrations for Case 2.

7 Examples showing lower bounds

This section presents two examples that give lower bounds on our results on MAP; each example is a well-structured instance of MAP. The first example gives a construction such that $opt \approx \frac{7}{4}\tau$. This shows that Theorem 6 is essentially tight. The second example gives a construction such that the cost of the solution computed by our algorithm is $\approx \frac{7}{4}opt$.

7.1 Optimal solution versus min-cost 2-edge cover

Proposition 28. *For any $k \in \mathbb{N}$, there exists a well-structured MAP instance G_k such that $\tau(G_k) \leq 4k + 3$ and $opt(G_k) \geq 7k + 3$.*

Proof. The graph G_k consists of a root 2-ec block B_0 and k copies J_1, \dots, J_k of a gadget subgraph J . The gadget subgraph J consists of 8 nodes v_1, \dots, v_8 and 11 edges; there are four zero-edges $v_1v_4, v_2v_3, v_5v_8, v_6v_7$, and seven unit-edges $v_1v_2, v_1v_7, v_2v_5, v_3v_4, v_3v_8, v_5v_6, v_7v_8$; see the subgraph induced by the nodes v_1, \dots, v_8 in Figure 14; observe that 8 of the 11 edges form two 4-cycles (namely, v_1, v_2, v_3, v_4, v_1 and v_5, v_6, v_7, v_8, v_5) and the other three edges are v_2v_5, v_3v_8 , and v_1v_7 .

Let B_0 be a 6-cycle w_1, \dots, w_6, w_1 that has 3 unit-edges and 3 zero-edges.




Figure 14: Example graph (with $k = 1$ copy of the 8-node gadget) where $\text{opt} \geq 7k + \epsilon$ and $\tau \leq 4k + \epsilon$, where ϵ is a constant. Edges of cost zero and one are illustrated by dashed and solid lines, respectively. B_0 is the root 2ec-block; ϵ is the cost of an optimal solution on B_0 .

G_k has two unit-edges between each copy of the gadget subgraph J_i ($i = 1, \dots, k$) and B_0 ; these two edges are incident to the nodes v_1 and v_3 of J_i (see the illustration in Figure 14) and to the nodes w_1 and w_4 of B_0 . Observe that the subgraph of G_k consisting of B_0 and the two 4-cycles (namely, v_1, v_2, v_3, v_4, v_1 and v_5, v_6, v_7, v_8, v_5) of each copy of the gadget subgraph is a (feasible) 2-edge cover of G_k of cost $4k + 3$. Hence, $\tau(G_k) \leq 4k + 3$.

Finally, we claim that $\text{opt}(G_k) \geq 7k + 3$. In what follows, we use OPT to denote an optimal solution of G_k , i.e., OPT denotes an arbitrary but fixed min-cost 2-ECSS of G_k . Clearly, OPT has to contain all the edges of B_0 as well as the two edges between B_0 and each copy of the gadget subgraph. Now, we focus on one copy J_i of the gadget subgraph, and let $\text{opt}(G_k, J_i)$ denote the cost of the edges of OPT incident to J_i . We will show that $\text{opt}(G_k, J_i) \geq 7$, hence, it follows that $\text{opt}(G_k) = 3 + \sum_{i=1}^k \text{opt}(G_k, J_i) \geq 7k + 3$. Since $\deg(v_4) = \deg(v_6) = 2$, OPT must pick the edges v_1v_4 and v_3v_4 as well as the edges v_5v_6 and v_6v_7 . Consider the cut $\delta_{G_k}(\{v_5, v_6, v_7, v_8\})$. This cut has three unit-edges: v_1v_7 , v_2v_5 , v_3v_8 . We have the following cases:

Case 1: OPT picks all three edges of the cut. Then $\text{opt}(G_k, J_i) \geq 7$.

Case 2: OPT does not pick all three edges of the cut. Then we will show that it must pick two edges from the cut and one more unit-edge, thus giving us $\text{opt}(G_k, J_i) \geq 7$. We have three subcases:

Case 2.1: $v_1v_7 \notin OPT$: then OPT must pick v_7v_8 and the other 2 edges of the cut.

Case 2.2: $v_2v_5 \notin OPT$: then OPT must pick v_1v_2 and the other 2 edges of the cut.

Case 2.3: $v_3v_8 \notin OPT$: then OPT must pick v_7v_8 and the other 2 edges of the cut.

Hence, $\text{opt}(G_k, J_i) \geq 7$ and we have $\text{opt}(G_k) \geq 7k + 3$. This completes the proof. \square

7.2 Optimal solution versus algorithm's solution

We present a family of graphs G_k , $k = 1, 2, 3, \dots$ for which the ratio of the cost of a solution obtained by our algorithm and the cost of an optimal solution approaches $\frac{7}{4}$. Let the solution subgraph found by applying our algorithm to G_k be denoted by $ALGO(G_k)$.

Proposition 29. *For any $k \in \mathbb{N}$, there exists a well-structured MAP instance G_k such that $\text{cost}(ALGO(G_k)) \geq 7k + 6$ and $\text{opt}(G_k) \leq 4k + 7$.*

Proof. Let J be a gadget on nodes $v_1, v_2, v_3, v_4, w_1, w_2, w_3, w_4$, with edges v_1v_4 , v_2v_3 , w_1w_2 , and w_3w_4 of cost zero, and edges v_1v_2 , v_3v_4 , w_1w_4 , w_2w_3 , $g = v_4w_1$ and $h = v_2v_4$ of cost one, as shown in Figure 15 (where dashed and solid lines represent edges of cost zero and one, respectively).

The graph $G_k = (V_k, E_k)$ is constructed as follows. We start with a 6-cycle $B_0 = b_1, b_2, b_3, b_4, b_5, b_6, b_1$ of unit-edges, of cost 6. We place k copies J_1, \dots, J_k of the gadget J in the following manner. Let v_j^i and w_ℓ^i denote the nodes v_j and w_ℓ of J_i , and let g_i and h_i denote the edges g and h of J_i .




Figure 15: The gadget J . Edges of cost zero and one are illustrated by dashed and solid lines, respectively.

First, we attach J_1 to B_0 by adding the two unit-edges $e_1 = b_1v_1^1$ and $f_1 = b_4v_3^1$. Then, for each $i \in \{2, \dots, k\}$, we attach J_i to J_{i-1} by adding the two unit-edges $e_i = w_2^{i-1}v_1^i$ and $f_i = w_3^{i-1}v_3^i$. The graph G_k is illustrated in Figure 16. Observe that G_k is a well-structured MAP instance.




Figure 16: The graph G_k has copies J_1, \dots, J_k of the gadget. All edges have cost one, except the zero-edges of J_1, \dots, J_k .

Note that the cost of any 2-edge cover is $\geq 4k + 6$, since it contains all edges of B_0 , as well as (at least) one unit-edge incident to each of the eight nodes of each gadget. W.l.o.g, $\widehat{D2}$ consists of B_0 and the two 4-cycles v_1, v_2, v_3, v_4, v_1 and w_1, w_2, w_3, w_4, w_1 of each gadget. Note that our choice of $\widehat{D2}$ is a bridgeless 2-edge cover.

Consider the working of the algorithm on G_k . Since $\widehat{D2}$ has no bridges, the algorithm proceeds to the gluing step. We use the notation of Section 6 to describe the working of the gluing step on G_k .

In each iteration of the gluing step, we choose the 2ec-block containing B_0 to be the root 2ec-block. In the first iteration, $R = B_0$ is the root 2ec-block, and the block B is the cycle v_1, v_2, v_3, v_4, v_1 of J_1 , i.e., the ear-augmentation step picks the “ear” R, B, R and takes the edges e, f (see Figure 12) to be the edges e_1, f_1 of G_k . Since the end nodes of e_1 and f_1 are non-adjacent in B , we apply case 2 of the gluing step by taking B' to be the 4-cycle w_1, w_2, w_3, w_4, w_1 of J_1 (see Figure 13b); moreover, we take \tilde{e} to be the edge g_1 (of G_k), and we take \tilde{P} to consist of h_1 and its two end nodes (in G_k). The algorithm buys the four edges e_1, f_1, g_1, h_1 and “sells” one unit-edge (say v_3v_4), so the algorithm incurs a cost of 7 for J_1 . In subsequent iterations, the same case of the gluing step is applied to each of the copies J_2, \dots, J_k of the gadget, hence, $ALGO(G_k)$ incurs a cost of 7 for each copy of the gadget; thus, we have $\text{cost}(ALGO(G_k)) = 7k + 6$.

On the other hand, the subgraph G^* of G_k (described below) is a 2-ECSS of cost $4k + 7$; $E(G^*)$ consists of the union of $k + 3$ sets of edges, namely, the set of edges $\{e_i, f_i, g_i, h_i\}$ for each $i \in \{1, \dots, k\}$, the set of zero-edges of G_k , $E(B_0)$, and the singleton $\{w_2^k w_3^k\}$ (the edge $w_2^k w_3^k$ is indicated by the right-most vertical line in Figure 16). Hence, $\text{opt}(G_k) \leq 4k + 7$. \square

Acknowledgments: We are grateful to several colleagues for their careful reading of preliminary drafts and for their comments. We thank an anonymous reviewer for a thorough review.

References

- [1] D. Adjiashvili. Beating approximation factor two for weighted tree augmentation with bounded costs. In: P. N. Klein (ed.) *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 2384–2399. SIAM, 2017.
- [2] R. Diestel. *Graph Theory (4th ed.)*. Graduate Texts in Mathematics, Volume 173. Springer-Verlag, Heidelberg, 2010.
- [3] G. Even, J. Feldman, G. Kortsarz, and Z. Nutov. A 1.8 approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2. *ACM Trans. Algorithms*, 5(2):21:1–17, 2009.
- [4] S. Fiorini, M. Groß, J. Könemann, and L. Sanità. Approximating weighted tree augmentation via Chvátal-Gomory cuts, In: A. Czumaj (ed.) *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 817–831. SIAM, 2018.
- [5] G. N. Frederickson and J. Jájá. Approximation algorithms for several graph augmentation problems. *SIAM J. Comput.*, 10(2):270–283, 1981.
- [6] G. N. Frederickson and J. Jájá. On the relationship between the biconnectivity augmentation and traveling salesman problems. *Theor. Comput. Sci.*, 19:189–201, 1982.
- [7] H. N. Gabow, M. X. Goemans, É. Tardos, and D. P. Williamson. Approximating the smallest k -edge connected spanning subgraph by LP-rounding. *Networks*, 53(4):345–357, 2009.
- [8] M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. *SIAM J. Comput.*, 24(2):296–317, 1995.
- [9] F. Grandoni, C. Kalaitzis, and R. Zenklusen. Improved approximation for tree augmentation: saving by rewiring. *Proc. 50th ACM Symposium on Theory of Computing, STOC*, pages 632–645, 2018.
- [10] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001.
- [11] M. Karpinski, M. Lampis, and R. Schmied. New inapproximability bounds for TSP. *J. Comput. Syst. Sci.*, 81(8):1665–1677, 2015.
- [12] G. Kortsarz and Z. Nutov. A simplified 1.5-approximation algorithm for augmenting edge-connectivity of a graph from 1 to 2. *ACM Trans. Algorithms*, 12(2):23:1–20, 2016.
- [13] L. C. Lau, R. Ravi, and M. Singh. *Iterative Methods in Combinatorial Optimization*. Cambridge Texts in Applied Mathematics (No. 46). Cambridge University Press, 2011.
- [14] L. Lovász and M. D. Plummer. *Matching Theory*, volume 367 of *AMS/Chelsea Publishing*. American Mathematical Society, 2009.
- [15] H. Nagamochi. An approximation for finding a smallest 2-edge connected subgraph containing a specified spanning tree. *Discrete Applied Mathematics*, 126:83–113, 2003.
- [16] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics, Volume 24. Springer-Verlag, Berlin Heidelberg, 2003.
- [17] A. Sebö and J. Vygen. Shorter tours by nicer ears: $7/5$ -approximation for the graph-TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. *Combinatorica*, 34(5):597–629, 2014.
- [18] S. Vempala and A. Vetta. Factor $4/3$ approximations for minimum 2-connected subgraphs. In K. Jansen and S. Khuller, (eds.) *Approximation Algorithms for Combinatorial Optimization, Third International Workshop, APPROX 2000, Proceedings*, LNCS 1913, pages 262–273. Springer, 2000.
- [19] H. Whitney. Non-separable and planar graphs. *Trans. Amer. Math. Soc.*, 34:339–362, 1932.
- [20] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.