

Round and Bipartize for Vertex Cover Approximation

Danish Kashaev ✉

Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

Guido Schäfer ✉

Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

ILLC, University of Amsterdam, The Netherlands

Abstract

The vertex cover problem is a fundamental and widely studied combinatorial optimization problem. It is known that its standard linear programming relaxation is integral for bipartite graphs and half-integral for general graphs. As a consequence, the natural rounding algorithm based on this relaxation computes an optimal solution for bipartite graphs and a 2-approximation for general graphs. This raises the question of whether one can interpolate the rounding curve of the standard linear programming relaxation in a beyond the worst-case manner, depending on how close the graph is to being bipartite. In this paper, we consider a round-and-bipartize algorithm that exploits the knowledge of an induced bipartite subgraph to attain improved approximation ratios. Equivalently, we suppose that we work with a pair (\mathcal{G}, S) , consisting of a graph with an odd cycle transversal.

If S is a stable set, we prove a tight approximation ratio of $1 + 1/\rho$, where $2\rho - 1$ denotes the odd girth (i.e., length of the shortest odd cycle) of the contracted graph $\tilde{\mathcal{G}} := \mathcal{G}/S$ and satisfies $\rho \in [2, \infty]$, with $\rho = \infty$ corresponding to the bipartite case. If S is an arbitrary set, we prove a tight approximation ratio of $(1 + 1/\rho)(1 - \alpha) + 2\alpha$, where $\alpha \in [0, 1]$ is a natural parameter measuring the quality of the set S . The technique used to prove tight improved approximation ratios relies on a structural analysis of the contracted graph $\tilde{\mathcal{G}}$, in combination with an understanding of the weight space where the fully half-integral solution is optimal. Tightness is shown by constructing classes of weight functions matching the obtained upper bounds. As a byproduct of the structural analysis, we also obtain improved tight bounds on the integrality gap and the fractional chromatic number of 3-colorable graphs. We also discuss algorithmic applications in order to find good odd cycle transversals, connecting to the MINUNCUT and COLOURING problems. Finally, we show that our analysis is optimal in the following sense: the worst case bounds for ρ and α , which are $\rho = 2$ and $\alpha = 1 - 4/n$, recover the integrality gap of $2 - 2/n$ of the standard linear programming relaxation, where n is the number of vertices of the graph.

2012 ACM Subject Classification Theory of computation → Rounding techniques

Keywords and phrases Combinatorial optimization, approximation algorithms, rounding algorithms, beyond the worst-case analysis

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2023.20

Category APPROX

Related Version *Full Version:* <https://arxiv.org/abs/2211.01699>

1 Introduction

In the vertex cover problem we are given a weighted graph $\mathcal{G} = (V, E, w)$, where $w : V \mapsto \mathbb{R}_+$ is a non-negative weight function on the vertices, and the goal is to find the minimal weight subset of vertices $U \subset V$ that covers every edge of the graph, i.e.,

$$\min\{w(U) \mid U \subset V, |U \cap \{i, j\}| \geq 1 \forall (i, j) \in E\}.$$



© Danish Kashaev and Guido Schäfer;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023).

Editors: Nicole Megow and Adam D. Smith; Article No. 20; pp. 20:1–20:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We denote by OPT an optimal subset of vertices for this problem, and by $w(OPT)$ the total weight of that solution. The vertex cover problem is known to be NP-complete [27] and APX-complete [38]. Moreover, it was shown to be NP-hard to approximate within a factor of $7/6$ in [23], a factor later improved to 1.36 in [16]. It is in fact NP-hard to approximate within a factor of $2 - \varepsilon$ for any fixed $\varepsilon > 0$ if the unique games conjecture is true [28].

A natural linear programming relaxation, as well as its dual, is given by:

$$\begin{array}{ll} \min \sum_{v \in V} w_v x_v & \max \sum_{e \in E} y_e \\ x_u + x_v \geq 1 \quad \forall (u, v) \in E & y(\delta(v)) \leq w_v \quad \forall v \in V \\ x_v \geq 0 \quad \forall v \in V & y_e \geq 0 \quad \forall e \in E \end{array}$$

For a given graph \mathcal{G} , we denote the primal linear program by $P(\mathcal{G})$ and the dual by $D(\mathcal{G})$. The integrality gap of the standard linear relaxation $P(\mathcal{G})$ on a graph of n vertices is upper bounded by $2 - 2/n$, a bound which is attained on the complete graph. In fact, a more fine-grained analysis shows that it is equal to $2 - 2/\chi^f(\mathcal{G})$, where $\chi^f(\mathcal{G})$ is the fractional chromatic number of the graph [45]. An integrality gap of $2 - \varepsilon$ is proved for a large class of linear programs in [4]. It is also known that any linear program which approximates vertex cover within a factor of $2 - \varepsilon$ requires super-polynomially many inequalities [10].

An important property of $P(\mathcal{G})$ is the fact that any extreme point solution x^* is half-integral, i.e., $x_v^* \in \{0, \frac{1}{2}, 1\}$ for any vertex $v \in V$ [35]. This gives rise to a straightforward rounding algorithm by solving $P(\mathcal{G})$ and outputting all vertices whose LP variable is at least a half, i.e., $U := \{v \in V \mid x_v^* \geq \frac{1}{2}\}$. It is easy to see that this is a 2-approximation, because $w(U) \leq 2w(OPT)$, see [24]. Moreover, it is known that $P(\mathcal{G})$ is integral for any bipartite graph [30]. As a consequence, the rounding algorithm returns an optimal solution if the graph is bipartite. This raises the question of whether we can interpolate the rounding curve of the standard linear program, depending on how close the graph is to being bipartite.

Set-up and algorithm

We consider the following setup. We are given a weighted non-bipartite graph $\mathcal{G} = (V, E, w)$ and an optimal solution $x^* \in \{0, \frac{1}{2}, 1\}$ to $P(\mathcal{G})$. We denote by $V_\alpha := \{v \in V \mid x_v^* = \alpha\}$ the vertices taking value α and by $\mathcal{G}_\alpha = \mathcal{G}[V_\alpha]$ the subgraph of \mathcal{G} induced by the vertices V_α for any $\alpha \in \{0, \frac{1}{2}, 1\}$. By a standard preprocessing step, we may assume that we only work on the graph $\mathcal{G}_{1/2}$, since any c -approximate solution on this reduced graph can be lifted to a c -approximate solution on the original graph by adding the nodes V_1 to the solution [35]. In addition, we suppose that we have knowledge of an odd cycle transversal S of $\mathcal{G}_{1/2}$, meaning that $\mathcal{G}_{1/2} \setminus S$ is a bipartite graph. Equivalently, S intersects every odd cycle of $\mathcal{G}_{1/2}$. The question of finding a good such odd cycle transversal is also tackled later in the paper.

We consider the following simple round and bipartize algorithm, detailed in Algorithm 1. It first solves $P(\mathcal{G})$, takes the vertices assigned value one by the linear program to the solution and removes all the integral nodes from the graph to arrive at $\mathcal{G}_{1/2}$. The algorithm then takes all the vertices in the set S to the solution, removes them from the graph and solves another (now integral) linear program to get the optimal solution on the bipartite remainder. These vertices are then also added to the solution.

The question studied is the following. What is the worst-case approximation ratio of the algorithm and which weight functions are attaining it? Our motivation to study this question comes from the structural difference between the polyhedron of $P(\mathcal{G})$ for bipartite and non-bipartite graphs. In particular, we are interested in identifying parameters of the

■ **Algorithm 1** Round and bipartize.

Input: Weighted graph $\mathcal{G} = (V, E, w)$, odd cycle transversal $S \subset V_{1/2}$

Output: Vertex cover $U \subset V$

- 1: Solve the linear program $P(\mathcal{G})$ to get $V_0, V_{1/2}$ and V_1
 - 2: Solve the integral linear program $P(\mathcal{G}_{1/2} \setminus S)$ to get $W \subset V_{1/2}$
 - 3: **return** $V_1 \cup S \cup W$
-

problem that enable us to derive more fine-grained bounds determining the approximation ratio of the algorithm, and allow to interpolate the rounding curve of the standard linear program from 1 to 2, depending on how far the graph is from being bipartite. As it turns out, the *odd girth*, i.e., the length of the shortest odd cycle, is a key parameter determining tight bounds on the approximation ratio. It is also a natural parameter, since a graph is bipartite if and only if it does contain an odd cycle. The larger the odd girth, the closer the graph is to being bipartite. It is also shown in [21] that graphs with a large odd girth admit a small cardinality odd cycle transversal.

Contributions and high-level view

We first do a pre-processing step and show that we may without loss of generality focus on weighted graphs $\mathcal{G} = (V, E, w)$ where the weights come from a certain weight space Q^W . Each edge has a dual weight $y_e \geq 0$ with a total sum of $y(E) = 1$, and the weight on each node is then determined by $w_v = y(\delta(v))$. This follows from the Nemhauser-Trotter theorem, complementary slackness and an appropriate normalization.

We then do the analysis under the assumption that S is a stable set, highlighting the main ideas of the analysis and the proof techniques. We show that the approximation ratio is upper bounded by $1 + 1/\rho$, where $2\rho - 1$ denotes the odd girth of the graph $\tilde{\mathcal{G}} := \mathcal{G}/S$, where all the vertices in S are contracted into a single node. Note that the parameter range is $\rho \in [2, \infty]$, with $\rho = \infty$ naturally corresponding to the case where $\tilde{\mathcal{G}}$ is bipartite. The proof technique involves a key concept, that we call *pairwise edge-separate* feasible vertex covers. Constructing k such covers allows to bound the approximation ratio by $1 + 1/k$. The construction of ρ such covers to get the result follows from a structural understanding of the contracted graph $\tilde{\mathcal{G}}$. As a byproduct, this structural understanding also allows to get improved bounds on the integrality gap and the fractional chromatic number of 3-colorable graphs. In particular, it even manages to compute an exact formula, depending on the odd girth, for the integrality gap and the fractional chromatic number of the contracted graph $\tilde{\mathcal{G}}$.

We then construct a class of weight functions $\mathcal{W} \subset Q^W$ where this upper bound holds with equality, thus showing that this proof technique obtains tight bounds and might have additional applications. This result can then be lifted to the case where S is a general set, by introducing an additional parameter α counting the total dual sum of the weights on the edges inside S , i.e. $\alpha = y(E[S])$. This then leads to an approximation ratio interpolating the rounding curve of the standard linear program with a tight bound of $(1 + 1/\rho)(1 - \alpha) + 2\alpha$ for any values of $\rho \in [2, \infty]$ and $\alpha \in [0, 1]$.

We then discuss algorithmic applications to find good such sets S , connecting to the MINUNCUT and COLOURING problems. Finally, we show that our analysis is optimal in the following sense: the worst case bounds for ρ and α , which are $\rho = 2$ and $\alpha = 1 - 4/n$, recover the integrality gap of $2 - 2/n$ of the standard linear programming relaxation for a graph on n vertices.

Implications and related work

Our analysis falls into the framework of beyond the worst-case analysis [41]. In particular, note that an odd cycle transversal always exists: we may simply take $S = V_{1/2}$, which recovers the standard 2-approximation algorithm for vertex cover. Depending on how S is chosen, our algorithm can however admit significantly better approximation ratios.

Our algorithm also connects to *learning-augmented algorithms*, which have access to some prediction in their input (e.g., obtained for instance through machine learning). This prediction is assumed to come without any worst-case guarantees, and the goal is then to take advantage of it by making the algorithm perform better when this prediction is good, while still keeping robust worst-case guarantees when it is off [7, 33, 39, 31, 2, 3]. In our case, assuming a prediction on the set S , robustness is guaranteed since we are never worse than a 2-approximation. In fact, even if the predicted set is not an odd cycle transversal, one may simply greedily add vertices to it until it becomes one, while still guaranteeing a 2-approximation. Otherwise, our results provide a precise understanding of how the approximation ratio improves depending on the predicted set S . In addition, once such a set S is found, the parameters α and ρ can easily be computed to see the improved guarantee on the approximation ratio. One may thus re-run the machine learning algorithm if the parameters give a bound very close to the worst-case of 2.

The odd cycle transversal number $\text{oct}(\mathcal{G})$ is defined as the minimum number of vertices needed to be removed in order to make the graph bipartite. The minimum odd cycle transversal problem has been studied in terms of fixed-parameter tractability [40, 29]. In particular, it is the first problem where the iterated compression technique has been applied [40], now a classical tool in the design of fixed-parameter tractable algorithms. The best known approximation algorithm for it has a ratio of $O(\sqrt{\log(n)})$ [1]. Another relevant concept is the *odd cycle packing number* $\text{ocp}(\mathcal{G})$, defined as the maximum number of vertex-disjoint odd cycles of \mathcal{G} and satisfying $\text{ocp}(\mathcal{G}) \leq \text{oct}(\mathcal{G})$. The related maximum stable set problem has been studied in terms of $\text{ocp}(\mathcal{G})$ in [11, 5, 15, 19].

A key property implying the integrality of a polyhedron is the *total unimodularity* (TU) of the constraint matrix describing the underlying problem, meaning that all the square subdeterminants of the matrix are required to lie in $\{-1, 0, 1\}$ (see for instance [43, 44]). In general we believe it is an interesting question to study whether one may exploit the knowledge of a TU substructure in an integer program to obtain improved approximation guarantees through rounding algorithms. In our case, the knowledge of an odd cycle transversal S is equivalent to the knowledge of an induced bipartite subgraph $\mathcal{G}' = \mathcal{G} \setminus S$, for which $P(\mathcal{G}')$ is TU. We hope the techniques introduced for the pair (\mathcal{G}, S) can help tackle other problems.

One technique which might also benefit from our analysis is iterative rounding, which requires solving a linear program at each iteration [32]. Having a better analysis for the case where the linear program becomes integral could potentially be used to reduce the number of iterations and allow for better guarantees, since iterative rounding can terminate at this step without losing solution quality.

Several different algorithms achieving approximation ratios of $2 - o(1)$ have been found for the weighted and unweighted versions of the vertex cover problem: [26, 22, 9, 34, 8, 25]. Another large body of work is interested in exact fixed parameter tractable algorithms for the decision version: [12, 6, 13, 14, 17, 36, 37, 46, 18].

2 Preliminaries

We define \mathbb{R}_+ to be the non-negative real numbers and $[k] := \{1, \dots, k\}$ to be the natural numbers up to $k \in \mathbb{N}$. For a vector $x \in \mathbb{R}^n$, we denote the sum of the coordinates on a subset by $x(A) := \sum_{i \in A} x_i$. A key property of $P(\mathcal{G})$ was introduced by Nemhauser and Trotter in [35]. It essentially allows to reduce an optimal solution $x^* \in \{0, \frac{1}{2}, 1\}^V$ to a fully half-integral solution by looking at the subgraph induced by the half-integral vertices. As before, we denote by $V_\alpha := \{v \in V \mid x_v^* = \alpha\}$ the vertices taking value α and by $\mathcal{G}_\alpha = \mathcal{G}[V_\alpha]$ the subgraph induced by the vertices V_α .

► **Theorem 1** (Nemhauser, Trotter [35]). *Let $x^* \in \{0, \frac{1}{2}, 1\}^V$ be an optimal extreme point solution to $P(\mathcal{G})$. Then, $w(\text{OPT}(\mathcal{G}_{1/2})) = w(\text{OPT}(\mathcal{G})) - w(V_1)$.*

► **Corollary 2.** *Let $x^* \in \{0, \frac{1}{2}, 1\}^V$ be an optimal solution to $P(\mathcal{G})$. If $S \subset V_{1/2}$ is a feasible vertex cover on $\mathcal{G}_{1/2}$ with approximation ratio at most ϕ , i.e., $w(S) \leq \phi w(\text{OPT}(\mathcal{G}_{1/2}))$, then $w(S) + w(V_1) \leq \phi w(\text{OPT}(\mathcal{G}))$.*

The previous corollary thus implies that we may restrict our attention to the graph $\mathcal{G}_{1/2}$, since any ϕ -approximate solution on this reduced instance can be lifted to a ϕ -approximate solution on the original graph by adding V_1 to the solution. Note that on the weighted graph $\mathcal{G}_{1/2}$, the solution $(\frac{1}{2}, \dots, \frac{1}{2})$ is optimal.

For a given set $S \subset V$, we define $\mathcal{G}' := \mathcal{G} \setminus S = (V', E')$ to be the graph obtained by removing the set S and all the incident edges to it. Hence, $E = E' \cup \delta(S) \cup E[S]$ where $\delta(S) = \{(u, v) \in E \mid u \in S, v \notin S\}$ and $E[S] := \{(u, v) \in E \mid u \in S, v \in S\}$. We also denote by $\tilde{\mathcal{G}} := \mathcal{G}/S = (\tilde{V}, \tilde{E})$ the graph obtained by contracting all the vertices in S into a single new node $v^S \in \tilde{V}$. We allow for multiple edges, but no self-loops. The only edges present in E but not in \tilde{E} are thus the ones with both endpoints in S , i.e., $E[S]$.

3 Weight Space

By Corollary 2, we may assume that every weighted graph \mathcal{G} we work with has the property that the fully half-integral solution $x = (\frac{1}{2}, \dots, \frac{1}{2})$ is an optimal solution to the linear program $P(\mathcal{G})$. In this section, we characterize the weight functions satisfying this assumption. The following lemma is a simple application of complementary slackness and the proof is omitted.

► **Lemma 3.** *Let $\mathcal{G} = (V, E)$ be a graph and let $w : V \rightarrow \mathbb{R}_+$ be a weight function. The feasible solution $(\frac{1}{2}, \dots, \frac{1}{2})$ to the linear program $P(\mathcal{G})$ is optimal if and only if there exists $y \in \mathbb{R}_+^E$ satisfying $y(\delta(v)) = w_v$ for every $v \in V$.*

Such instances have been called *edge-induced* in [15, 19], in the sense that the dual values on the edges are free parameters, and the weights on the nodes are determined once the dual values are fixed. Such instances also satisfy:

$$w(V) = \sum_{v \in V} w_v = \sum_{v \in V} y(\delta(v)) = \sum_{v \in V} \sum_{e \in E} y_e \mathbf{1}_{\{e \in \delta(v)\}} = \sum_{e \in E} y_e \sum_{v \in V} \mathbf{1}_{\{e \in \delta(v)\}} = 2 y(E).$$

Observe that the approximation ratio of a feasible solution $U \subset V$ is defined as $w(U)/w(\text{OPT}(\mathcal{G}))$ and is invariant under scaling of the weights. We thus make a normalization ensuring that the optimal LP solution has objective value one, i.e., $w(V)/2 = y(E) = 1$, to get the following weight space polytope:

$$Q^W := \{w \in \mathbb{R}_+^V \mid \exists y \in [0, 1]^E \text{ such that } y(E) = 1 \text{ and } w_v = y(\delta(v)) \quad \forall v \in V\}.$$

We end this section by showing that this normalization of the weight space allows us to get a convenient lower bound on $w(OPT(\mathcal{G}))$.

► **Lemma 4.** *Let $\mathcal{G} = (V, E)$ be a graph. For any $w \in Q^W$, $w(OPT(\mathcal{G})) \geq 1$.*

Proof. Since $w \in Q^W$, we know that the fully half-integral solution is an optimal linear programming solution, showing $1 = w(V)/2 \leq w(OPT(\mathcal{G}))$, by feasibility of $OPT(\mathcal{G})$. ◀

4 Round and Bipartize

4.1 Algorithm

This section is devoted to the analysis of the approximation ratio of the algorithm and is the main contribution of the paper. We assume that we are given as input a pair (\mathcal{G}, S) consisting of a weighted graph and an odd cycle transversal $S \subset V$. By Corollary 2, we may assume that the weight function satisfies $w \in Q^W$. By the previous section, there are dual edge weights $y_e \geq 0$ such that $w_v = y(\delta(v))$ for every $v \in V$ and which satisfy $\sum_{e \in E} y_e = 1$.

The algorithm is now very simple. First, take the vertices in $S \subset V$ to the cover and remove them from the graph. Then solve the integral linear program $P(\mathcal{G} \setminus S)$ and take the vertices having LP value one to the cover. The *approximation ratio*, given a weight function w , is thus defined as

$$R(w) := \frac{w(S) + w(OPT(\mathcal{G} \setminus S))}{w(OPT(\mathcal{G}))}. \quad (1)$$

For simplicity of notation, we omit the dependence on w of $OPT(\mathcal{G})$ and $OPT(\mathcal{G} \setminus S)$. As a reminder, the bipartite graph $\mathcal{G} \setminus S$ is denoted by $\mathcal{G}' = (V', E')$. The vertex contracted graph \mathcal{G}/S is denoted by $\tilde{\mathcal{G}} = (\tilde{V}, \tilde{E})$ and the contracted node is denoted by v_S .

4.2 Stable Set to Bipartite

We assume in this section that S is a stable set. We will then generalize the results obtained here in a natural way to the most general setting of an arbitrary set S . We now state our main theorem of this section.

► **Theorem 5.** *Let (\mathcal{G}, S) be the input to the rounding algorithm, with S being a stable set. For any $w \in Q^W$, the approximation ratio satisfies $R(w) \leq 1 + 1/\rho$, where $2\rho - 1$ is the odd girth of the contracted graph $\tilde{\mathcal{G}}$ and satisfies $\rho \in [2, \infty]$. Moreover, this bound is tight and is attained for a class of weight functions $\mathcal{W} \subset Q^W$.*

► **Remark 6.** We define the odd girth of a bipartite graph as being ∞ .

► **Definition 7.** *Let (\mathcal{G}, S) be a pair consisting of a graph with an odd cycle transversal S . For a feasible vertex cover $U \subset V \setminus S$ of the bipartite graph $\mathcal{G}' = \mathcal{G} \setminus S$, we define*

$$E_U := \{(u, v) \in E \mid u \in U, v \in U \text{ or } u \in U, v \in S\}.$$

In words, these are the edges with either both endpoints in the cover U , or with one endpoint in U and one in S .

► **Definition 8.** *Let (\mathcal{G}, S) be a pair consisting of a graph with an odd cycle transversal S . Feasible vertex covers U_1, \dots, U_k of the bipartite graph $\mathcal{G}' = \mathcal{G} \setminus S$ are defined to be pairwise edge-separate if the edge sets $\{E_{U_1}, \dots, E_{U_k}\}$ are pairwise disjoint.*

► **Remark 9.** We will often say that covers are pairwise edge-separate for the pair (\mathcal{G}, S) . It is however worth emphasizing that these covers are defined on the bipartite graph $\mathcal{G}' = \mathcal{G} \setminus S$.

This definition turns out to be the key concept for us in order to prove improved bounds on the approximation ratio of the algorithm, as shown by the next lemma.

► **Lemma 10.** *Let (\mathcal{G}, S) be the input to the rounding algorithm, with S being a stable set. If there exists k pairwise edge-separate feasible vertex covers of the bipartite graph $\mathcal{G}' = \mathcal{G} \setminus S$, then, for every $w \in Q^W$, the approximation ratio of the algorithm satisfies $R(w) \leq 1 + 1/k$.*

Proof. Let $w \in Q^W$ and let $y \in \mathbb{R}_+^E$ be the corresponding dual solution satisfying $w_v = y(\delta(v))$ and $y(E) = 1$. We denote by $\{U_1, \dots, U_k\}$ the pairwise edge-separate covers of $\mathcal{G}' = (V', E')$. We can now write down the weights of S and every feasible cover U_i with the help of the dual variables:

$$\begin{aligned} w(S) &= \sum_{v \in S} w_v = \sum_{v \in S} y(\delta(v)) = y(\delta(S)) \\ w(U_i) &= \sum_{v \in U_i} w_v = \sum_{v \in U_i} y(\delta(v)) = y(E') + y(E_{U_i}) \quad \forall i \in [k] \end{aligned}$$

The first equality holds because S is a stable set and thus only has edges crossing the set. The second equality holds because every U_i counts the dual value y_e of every $e \in E'$ at least once, by feasibility of the cover, and thus giving a contribution of $y(E')$. The edges in E_{U_i} then give an additional contribution of $y(E_{U_i})$.

By Lemma 4, the approximation ratio satisfies:

$$\begin{aligned} R(w) &= \frac{w(S) + w(OPT(\mathcal{G} \setminus S))}{w(OPT(\mathcal{G}))} \leq w(S) + w(OPT(\mathcal{G}')) \leq w(S) + \min_{i \in [k]} w(U_i) \\ &= y(\delta(S)) + y(E') + \min_{i \in [k]} y(E_{U_i}) = 1 + \min_{i \in [k]} y(E_{U_i}) \leq 1 + \frac{1}{k}. \end{aligned}$$

The last equality follows from the fact that $E = E' \cup \delta(S)$ and $y(E) = 1$. The last inequality follows from the fact that the edge sets $\{E_{U_i}\}_{i \in [k]}$ are pairwise disjoint and have a dual sum of at most one, since the total sum of the edges of the graph is $y(E) = 1$. This minimum can thus be upper bounded by $1/k$. ◀


In order to prove the upper bound in Theorem 5, we thus need to construct ρ pairwise edge-separate covers of $\mathcal{G}' = \mathcal{G} \setminus S$. The key for being able to do that is to analyze the structure of the contracted graph $\tilde{\mathcal{G}} = \mathcal{G}/S$, where S is contracted into a single node v_S .

► **Lemma 11.** *Let (\mathcal{G}, S) be a graph with an odd cycle transversal S . If the contracted graph $\tilde{\mathcal{G}}$ contains an odd cycle, then there exists ρ edge-separate feasible covers for the pair $(\tilde{\mathcal{G}}, v_S)$, where $2\rho - 1$ is the odd girth of $\tilde{\mathcal{G}}$.*

Proof. We now dive deeper into the structure of the bipartite graph $\mathcal{G} \setminus S = \tilde{\mathcal{G}} \setminus v_S$. By assumption, this graph admits a bipartition $A \cup B$ of the vertices. Let us assume that it has k connected components $A_1 \cup B_1, \dots, A_k \cup B_k$, all of which are bipartite as well, where $A = \bigcup_i A_i$ and $B = \bigcup_i B_i$. We now fix an arbitrary such component $A_j \cup B_j$.

- If v_S has an incident edge to both A_j and B_j , then this component contains (if including v_S) an odd cycle of $\tilde{\mathcal{G}}$. This holds since any path between a node in A_j and a node in B_j has odd length.
- If v_S has incident edges with only one side, we assume without loss of generality that this side is A_j . One could simply switch both sides in the other case while still keeping a valid bipartition of the graph $\tilde{\mathcal{G}} \setminus v_S$.

20:8 Round and Bipartize for Vertex Cover Approximation



■ **Figure 1** The layers of a bipartite graph $\tilde{\mathcal{G}} \setminus v_S = (A \cup B, E')$ with $\rho = 4$. The blue square vertices correspond to $N(v_S)$, where the two left ones are $\mathcal{L}_0 = N_A(v_S)$ and the two right ones are $N_B(v_S)$.

- If v_S does not have incident edges with either of the two sides, then $A_j \cup B_j$ is a connected component of $\tilde{\mathcal{G}}$. We call such components *dummy* components and denote by $A_d \cup B_d$ the bipartite graph formed by taking the union of all the dummy components.

We denote $N_A(v_S) = N(v_S) \cap A$ and $N_B(v_S) = N(v_S) \cap B$. We now split the graph into layers, where each layer corresponds to the nodes at the same shortest path distance from $N_A(v_S)$. More precisely, we define

$$\mathcal{L}_i := \left\{ v \in A \cup B \mid d(N_A(v_S), v) = i \right\} \quad \text{for } i \in \{0, \dots, q\} \quad (2)$$

where $d(N_A(v_S), v)$ represents the unweighted shortest path distance between v and a vertex in $N_A(v_S)$. The parameter q is defined to be the maximal finite distance from $N_A(v_S)$ in the graph $\tilde{\mathcal{G}}$. An important observation is the fact that these layers are alternately included in one side of the bipartition, see Figure 1 for an illustration of the construction.

If the graph $\tilde{\mathcal{G}}$ is not connected, note that $d(N_A(v_S), v) = \infty$ for the vertices v lying in dummy components. In order to add the dummy components to the layers and keep alternation between the two sides of the bipartition, we define the last two layers to either be $\{\mathcal{L}_{q+1} := A_d, \mathcal{L}_{q+2} := B_d\}$ or $\{\mathcal{L}_{q+1} := B_d, \mathcal{L}_{q+2} := A_d\}$, depending on which side of the bipartition the last connected layer \mathcal{L}_q lies. We now have that $\mathcal{L}_i \subset A$ if i is even, and $\mathcal{L}_i \subset B$ if i is odd. In fact,

$$A = \bigcup_{i=0}^{\lfloor l/2 \rfloor} \mathcal{L}_{2i} \quad \text{and} \quad B = \bigcup_{i=1}^{\lceil l/2 \rceil} \mathcal{L}_{2i-1},$$

where the parameter $l \in \mathbb{N}$ represents the index of the last layer: if $\tilde{\mathcal{G}}$ is connected, then $l = q$, otherwise $l = q + 2$. Notice also that $\mathcal{L}_0 = N_A(v_S)$. However, $N_B(v_S)$ may now have several different vertices in different layers, see Figure 1.


Let $C \subset V$ be an arbitrary odd cycle of $\tilde{\mathcal{G}}$. Notice that this cycle contains v_S , a vertex from $N_A(v_S)$ and a vertex from $N_B(v_S)$, since $\tilde{\mathcal{G}} \setminus v_S$ is bipartite and therefore does not contain an odd cycle. Any odd cycle C in $\tilde{\mathcal{G}}$ thus corresponds to an odd path between a vertex in $N_A(v_S) = \mathcal{L}_0$ and a vertex in $N_B(v_S)$. By the assumption that the shortest odd cycle length of $\tilde{\mathcal{G}}$ is $2\rho - 1$, the first layer having a non-empty intersection with $N_B(v_S)$ is $\mathcal{L}_{2\rho-3}$. A shortest odd cycle of length $2\rho - 1$ therefore corresponds to an odd path of length $2\rho - 3$ between \mathcal{L}_0 and a vertex in $\mathcal{L}_{2\rho-3} \cap N_B(v_S)$, see Figure 1 for an illustration. We now define edges connecting two consecutive layers \mathcal{L}_i and \mathcal{L}_{i+1} as follows:

$$E[\mathcal{L}_i, \mathcal{L}_{i+1}] := \{(u, v) \in E' \mid u \in \mathcal{L}_i, v \in \mathcal{L}_{i+1}\} \quad \forall i \in \{0, \dots, l-1\}.$$

We also denote by

$$\delta_A(v_S) = \{(v_S, u) \in \tilde{E} \mid u \in A\}, \quad \delta_B(v_S) = \{(v_S, u) \in \tilde{E} \mid u \in B\}$$

the incident edges to v_S respectively connecting to A and B .



■ **Figure 2** The ρ feasible covers of \mathcal{G}' constructed in the proof of Lemma 11.

We are now ready to construct our desired ρ pairwise edge-separate covers of $\tilde{\mathcal{G}} \setminus v_S$, that we denote by U_1, \dots, U_ρ and illustrated in Figure 2. Firstly, notice that taking one side of the bipartition is a feasible vertex cover. We thus define $U_\rho = A$ and $U_{\rho-1} = B$. Observe that we then have $E_{U_\rho} = \delta_A(v_S)$ and $E_{U_{\rho-1}} = \delta_B(v_S)$. We now construct $\rho - 2$ additional covers with the help of the layers. If $\rho \neq 2$, fix a $j \in [\rho - 2]$, and start the cover U_j by taking the two consecutive layers \mathcal{L}_{2j-1} and \mathcal{L}_{2j} . Complete this cover by taking remaining layers alternatingly (hence always skipping one) until covering every edge of the graph. Notice that this cover has an empty intersection with $N(v_S)$. We then have that

$$E_{U_\rho} = \delta_A(v_S), \quad E_{U_{\rho-1}} = \delta_B(v_S), \quad E_{U_j} = E[\mathcal{L}_{2j-1}, \mathcal{L}_{2j}] \quad \forall j \in [\rho - 2],$$

which are all pairwise disjoint edge sets, finishing the proof. ◀

We now have all the tools to prove the upper bound of Theorem 5.


Proof. Assume first that $\rho < \infty$, meaning that $\tilde{\mathcal{G}}$ contains an odd cycle. By Lemma 11, there exists ρ pairwise edge-separate covers for the pair $(\tilde{\mathcal{G}}, v_S)$. These covers are then still edge-separate for the pair (\mathcal{G}, S) , since the bipartite graph is the same in both cases, i.e. $\mathcal{G}' = \tilde{\mathcal{G}} \setminus v_S = \mathcal{G} \setminus S$. This finishes the proof by Lemma 10.

If $\rho = \infty$, then $\tilde{\mathcal{G}}$ is bipartite, with a bipartition $\tilde{A} \cup \tilde{B}$. Assume without loss of generality that $v^S \in \tilde{A}$. Note that $\tilde{E} = E' \cup \delta(v^S)$ and thus $1 = y(\tilde{E}) = y(E') + y(\delta(v^S))$. Any feasible cover of $\mathcal{G}' = \mathcal{G} \setminus S$ needs to count the dual value of every edge in E' at least once. Taking the cover $\tilde{A} \setminus v^S$ counts every edge in E' exactly once, showing that $w(OPT(\mathcal{G} \setminus S)) = y(E')$. Hence, $R(w) \leq w(S) + w(OPT(\mathcal{G} \setminus S)) = y(\delta(v_S)) + y(E') = 1$. ◀

We now show that this bound is tight and is attained for a class of weight functions $w \in \mathcal{W}$ for any such graph \mathcal{G} and stable set S . For the case where $\rho = \infty$, it is clear that the approximation ratio always satisfies $R(w) \geq 1$, showing that the bound in Theorem 5 is tight for any $w \in Q^W$. We thus assume that $\rho < \infty$. Let \mathcal{C} be all the shortest odd cycles (of length $2\rho - 1$) of the graph $\tilde{\mathcal{G}}$, each of which is containing v_S . For every such cycle $C \in \mathcal{C}$, we define the following dual function on the edges $y^C : \tilde{E} \rightarrow \mathbb{R}_+$: set both dual edges incident to v_S to $1/\rho$ and then alternatingly set the dual edges to 0 and $1/\rho$ along the odd cycle. For any edge outside of C , set its dual value to 0. Such a solution clearly satisfies $y^C(\tilde{E}) = 1$. We now take the convex hull of all these functions:

$$\mathcal{Y} := \left\{ y : \tilde{E} \rightarrow \mathbb{R}_+ \mid y = \sum_{C \in \mathcal{C}} \lambda^C y^C, \quad \sum_{C \in \mathcal{C}} \lambda^C = 1, \quad \lambda^C \geq 0 \quad \forall C \in \mathcal{C} \right\}.$$

20:10 Round and Bipartize for Vertex Cover Approximation



■ **Figure 3** An example of a weight function $w \in \mathcal{W}$ obtained by a convex combination of two basic weight functions of shortest odd cycles.

Because of the one-to-one correspondence between the edge sets \tilde{E} and E , due to the fact that S is a stable set, we can naturally define a weight function on the original vertex set once we fix a $y \in \mathcal{Y}$ by setting $w_v := y(\delta(v))$ for every $v \in V$. We define the space of all such weight functions as $\mathcal{W} := \{w : V \rightarrow \mathbb{R}_+ \mid w_v = y(\delta(v)) \forall v \in V, y \in \mathcal{Y}\}$.

► **Theorem 12.** *For any weight function $w \in \mathcal{W}$, the approximation ratio satisfies $R(w) = 1 + 1/\rho$, where $2\rho - 1$ is the odd girth of $\tilde{\mathcal{G}}$ and satisfies $\rho \in [2, \infty)$.*

Proof. Let \mathcal{C} be the set of all the shortest odd cycles (of length $2\rho - 1$) of the graph $\tilde{\mathcal{G}}$ and let $w \in \mathcal{W}$ with the corresponding $y = \sum_{C \in \mathcal{C}} \lambda^C y^C$. Notice that, for any subset of vertices $U \subset V'$ of the bipartite graph \mathcal{G}' , we can count its weight as

$$\begin{aligned} w(U) &= \sum_{v \in U} w_v = \sum_{v \in U} y(\delta(v)) = \sum_{v \in U} \sum_{C \in \mathcal{C}} \lambda^C y^C(\delta(v)) \\ &= \sum_{v \in U} \sum_{C \in \mathcal{C}} \frac{\lambda^C}{\rho} 1_{\{v \in C\}} = \frac{1}{\rho} \sum_{C \in \mathcal{C}} \lambda^C \sum_{v \in U} 1_{\{v \in C\}} = \frac{1}{\rho} \sum_{C \in \mathcal{C}} \lambda^C |U \cap C|. \end{aligned} \quad (3)$$

The end of the proof now heavily uses the decomposition of $\tilde{\mathcal{G}}$ into layers described in (2). Notice that every odd cycle $C \in \mathcal{C}$ intersects each layer \mathcal{L}_i for $i \in \{0, \dots, 2\rho - 3\}$ exactly once. Therefore, by (3), $w(\mathcal{L}_i) = 1/\rho$ for every $i \in \{0, \dots, 2\rho - 3\}$. We now claim that

$$w(OPT(\mathcal{G})) = 1, \quad w(OPT(\mathcal{G}')) = \frac{\rho - 1}{\rho} \quad \text{and} \quad w(S) = \frac{2}{\rho}.$$


The fact that $w(OPT(\mathcal{G})) \geq 1$ follows from Lemma 4. For the reverse inequality, notice that it is possible to take a feasible cover by taking exactly ρ layers in addition to the zero weight vertices, for instance $\mathcal{L}_0 \cup \mathcal{L}_2 \cup \mathcal{L}_3 \cup \mathcal{L}_5 \cdots \cup \mathcal{L}_{2\rho-3}$, showing $w(OPT(\mathcal{G})) \leq 1$.

Observe now that $w(OPT(\mathcal{G}')) = w(OPT(\mathcal{G} \setminus S)) = w(OPT(\tilde{\mathcal{G}} \setminus v_S))$. After removal of v_S , every cycle $C \in \mathcal{C}$ becomes a path of length $2\rho - 3$ (and thus consisting of $2\rho - 2$ vertices), with one vertex in each layer \mathcal{L}_i for $i \in \{0, \dots, 2\rho - 3\}$. By feasibility, $OPT(\mathcal{G}')$ has to contain at least $\rho - 1$ vertices for every such path. Using (3), we infer $w(OPT(\mathcal{G}')) \geq (\rho - 1)/\rho$. For the reverse inequality, taking $\rho - 1$ layers alternatively, such as $\mathcal{L}_0 \cup \mathcal{L}_2 \cup \mathcal{L}_4 \cdots \cup \mathcal{L}_{2\rho-4}$, as well as the zero weight vertices, builds a feasible cover of weight exactly $(\rho - 1)/\rho$.

Finally, notice that $w(S) = w(v_S) = 2/\rho$ because every $C \in \mathcal{C}$ contains v_S . By combining the three equalities, we get the desired result $R(w) = 1 + 1/\rho$. ◀

Integrality gap and fractional chromatic number

Observe that the graphs we work with in this section have chromatic number $\chi(\mathcal{G}) = 3$. We focus now on the integrality gap and show again that the odd girth plays a key role. A result that we use here is given by Singh in [45], which relates the integrality gap with the fractional chromatic number of a graph: $IG(\mathcal{G}) = 2 - 2/\chi^f(\mathcal{G})$.



■ **Figure 4** The plot of the approximation ratio $R(w)$ with respect to the parameter $\alpha \in [0, 1]$.

► **Theorem 13.** Let \mathcal{G} be a 3-colorable graph with color classes $V = V_1 \cup V_2 \cup V_3$.

$$\chi^f(\mathcal{G}) \leq 2 + \min_{i \in \{1,2,3\}} \frac{1}{\rho_i - 1}, \quad IG(\mathcal{G}) \leq 1 + \min_{i \in \{1,2,3\}} \frac{1}{2\rho_i - 1}$$

where $2\rho_i - 1$ is the odd girth of the contracted graph \mathcal{G}/V_i for each $i \in \{1, 2, 3\}$. Moreover, equality holds if one color class only contains one vertex.

In particular, we manage to compute exact formulas of

$$\chi^f(\tilde{\mathcal{G}}) = 2 + \frac{1}{\rho - 1}, \quad IG(\tilde{\mathcal{G}}) = 1 + \frac{1}{2\rho - 1}$$

for the contracted graph $\tilde{\mathcal{G}}$. These statements generalize a result shown for the (odd) cycle graph in [4, 42]. The proofs are left to the appendix due to space constraints and heavily use the structural decomposition of the contracted graph $\tilde{\mathcal{G}}$ into the layers (2), implying that it may have further applications.

4.3 Arbitrary Set to Bipartite

We now consider the setting where S is now an arbitrary set. Our guarantee on the approximation ratio will now also depend on the total sum of the dual variables on the edges inside of the set S . We denote this sum by $\alpha := y(E[S]) \in [0, 1]$.

► **Theorem 14.** For any $w \in Q^W$, the approximation ratio satisfies

$$R(w) \leq \left(1 + \frac{1}{\rho}\right) (1 - \alpha) + 2\alpha \quad \text{with } \alpha \in [0, 1] \text{ and } \rho \in [2, \infty]$$

where $2\rho - 1$ denotes the odd girth of the contracted graph $\tilde{\mathcal{G}}$. Moreover, these bounds are tight and are attained for any $\alpha \in [0, 1]$ and any $\rho \in [2, \infty]$.

Proof. We only prove here the upper bound with $\rho < \infty$ and leave the remaining statements to the appendix. The proof essentially follows the same arguments as the one of Lemma 10 with the α parameter incorporated, and we thus only highlight the main differences. We decompose the weight of the set S with respect to the dual variables. The edges in $E[S]$ are counted twice, whereas the edges in $\delta(S)$ are counted once:

$$w(S) = 2\alpha + y(\delta(S)). \tag{4}$$

20:12 Round and Bipartize for Vertex Cover Approximation

Consider the contracted graph $\tilde{\mathcal{G}} = \mathcal{G}/S = (\tilde{V}, \tilde{E})$ and denote by v^S the contracted node. The edge set of this graph is now $\tilde{E} = \delta(S) \cup E'$, since the edges in $E[S]$ have been collapsed. By Lemma 11, we can construct ρ edge-separate covers $U_1, \dots, U_\rho \subset \tilde{V} \setminus v^S$ for the pair $(\tilde{\mathcal{G}}, v^S)$. These covers are still edge-separate for the pair (\mathcal{G}, S) , implying

$$w(\text{OPT}(\mathcal{G} \setminus S)) \leq \min_{i \in [\rho]} w(U_i) = y(E') + \min_{i \in [\rho]} y(E_{U_i}) \leq y(E') + \frac{1 - \alpha}{\rho}. \quad (5)$$

The first inequality holds since every U_i is a feasible cover of $\mathcal{G} \setminus S$. The second equality holds by counting the weight of a cover U_i in terms of the dual edges. The last inequality holds because the edge sets $\{E_{U_i}\}_{i \in [\rho]}$ are pairwise disjoint, and their total dual sum is at most $1 - \alpha$. Combining Lemma 4, (4) and (5),

$$R(w) \leq 2\alpha + y(\delta(S)) + y(E') + \frac{1 - \alpha}{\rho} = 1 + \alpha + \frac{1 - \alpha}{\rho} = \left(1 + \frac{1}{\rho}\right)(1 - \alpha) + 2\alpha. \quad \blacktriangleleft$$

5 Algorithmic applications

We focus in this section on efficient ways to find odd cycle transversals with a low value for the α parameter. In fact, once such a set S is found, there can also be freedom in the choice of the dual solution in order to optimize the α parameter. This motivates the following definition.

► **Definition 15.** Let (\mathcal{G}, S, y, w) be a graph with an odd cycle transversal $S \subset V$, weights $w \in Q^W$ and a dual solution $y \in \mathbb{R}_+^E$. A tuple $(\mathcal{G}', S', y', w')$ is approximation preserving if

$$w(S) + w(\text{OPT}(\mathcal{G} \setminus S)) \leq w'(S') + w'(\text{OPT}(\mathcal{G}' \setminus S')).$$

Moreover, we say that $\alpha \in [0, 1]$ is valid for the pair (\mathcal{G}, S) if there exists an approximation preserving $(\mathcal{G}', S', y', w')$ such that $\alpha = y'(E[S'])$.

Finding a valid $\alpha \in [0, 1]$ would directly allow us to use it in the bound of Theorem 14, where the ρ parameter would correspond to the one of the approximation preserving graph. We present here an application if a coloring of a graph can be found efficiently.

► **Theorem 16.** Let $\mathcal{G} = (V, E)$ be a graph with weights $w \in Q^W$ that can be k -colored in polynomial time for $k \geq 4$. There exists an efficiently findable set $S \subset V$ bipartizing the graph and a valid α such that $\alpha \leq 1 - 4/k$, leading to an approximation ratio of

$$R(w) \leq 2 - \frac{4}{k} \left(1 - \frac{1}{\rho}\right)$$

Proof. Let us denote by V_1, \dots, V_k the k independent sets defining the color classes of the graph \mathcal{G} . We assume without loss of generality that they are ordered by weight $w(V_1) \leq w(V_2) \leq \dots \leq w(V_k)$. Since $w(V) = 2$, the two color classes with the largest weights satisfy $w(V_{k-1}) + w(V_k) \geq 4/k$. We define the bipartizing set to be the remaining color classes: $S := V_1 \cup \dots \cup V_{k-2}$. We denote by $y \in \mathbb{R}_+^E$ the dual solution satisfying complementary slackness and $y(E) = 1$.

We now define an approximation preserving $(\mathcal{G}', S', y', w')$ in the following way. Let $\mathcal{G}' = K_k$ be the complete graph on k vertices, denoted by $\{v_1, \dots, v_k\}$. The weights are defined to be

$$w'(v_i) := w(V_i) \quad \text{and} \quad y'(v_i, v_j) := y(E[V_i, V_j])$$

for every $i, j \in [k]$. These clearly satisfy the complementary slackness condition $y'(\delta(v_i)) = w'(v_i)$ for every $i \in [k]$. The bipartizing set is defined to be $S' := \{v_1, \dots, v_{k-2}\}$. This tuple is approximation preserving since $w(S) = w'(S')$ and $w(OPT(\mathcal{G} \setminus S)) \leq w'(OPT(\mathcal{G}' \setminus S'))$. In order to prove the theorem, we still need to tweak the dual solution y' to ensure $\alpha := y'(E[S']) \leq 1 - 4/k$. Observe that $w'(v_{k-1}) + w'(v_k) \geq 4/k$.

- (i) If $y'(v_{k-1}, v_k) = 0$, then the result follows since in that case $y'(\delta(S')) \geq 4/k$ and thus $y'(E[S']) \leq 1 - 4/k$.
- (ii) If $y'(E[S']) = 0$, then the result trivially follows as well.

Suppose thus that $y'(E[S']) > 0$ and $y'(v_{k-1}, v_k) > 0$. Pick an arbitrary edge $(v_i, v_j) \in E[S']$ satisfying $y'(v_i, v_j) > 0$ and consider the 4-cycle (v_i, v_j, v_{k-1}, v_k) . Notice that alternatively increasing and decreasing the dual values on the edges of this cycle by a small amount $\epsilon > 0$ gives another feasible dual solution satisfying the complementary slackness condition. More formally, we set $\epsilon := \min\{y'(v_i, v_j), y'(v_{k-1}, v_k)\}$, decrease $y'(v_i, v_j)$ and $y'(v_{k-1}, v_k)$ by ϵ , while increasing $y'(v_j, v_{k-1})$ and $y'(v_k, v_i)$ by the same amount. Observe that this leads to either (v_i, v_j) or (v_{k-1}, v_k) dropping to dual value zero. We can repeat this procedure until either $y'(E[S']) = 0$ or $y'(v_{k-1}, v_k) = 0$, finishing the proof of the theorem. ◀

We now claim that this result is optimal in the following sense. Consider an n -vertex graph. It is known that the integrality gap of the standard linear programming relaxation for vertex cover is upper bounded by $2 - 2/n$, a bound which is attained on the complete graph. This implies that any approximation algorithm lower bounding $w(OPT)$ by comparing it to the optimal LP solution, as we do in Lemma 4, cannot do better than $2 - 2/n$ in the worst case. Setting $\rho = 2$ in Theorem 16, which corresponds to the worst case since $\rho \in [2, \infty]$, recovers this bound and a result of Hochbaum in [25].

We now make another connection with the MINUNCUT problem, which is defined on a graph $\mathcal{G} = (V, E)$ with weights y_e for every edge $e \in E$. The goal of the problem is to find a cut of the graph which minimizes the total weight of uncut edges. This problem is NP-hard and admits a $O(\sqrt{\log(n)})$ approximation [1]. We call a MINUNCUT instance *light* if its optimal solution is bounded above by $y(E)/\Omega(\log(n))$.

▶ **Theorem 17.** *For any light MINUNCUT instance, combining the $O(\sqrt{\log(n)})$ approximation in [1] with Algorithm 1 outputs a vertex cover with approximation ratio at most $R(w) \leq 1 + 1/\rho + o(1)$, where $2\rho - 1$ is the odd girth of the contracted graph $\tilde{\mathcal{G}}$ and satisfies $\rho \in [2, \infty]$.*

Proof. The key observation is that any feasible solution to the MINUNCUT problem of value α gives an odd cycle transversal S satisfying $y(E[S]) = \alpha$. Indeed, let $(C, V \setminus C)$ be a feasible solution, where the total weight of uncut edges is α . Observe that removing all the uncut edges makes C and $V \setminus C$ become stable sets, implying that the remaining graph is bipartite. We then define $S \subset V$ to be all the nodes incident to the uncut edges. Since removing all the nodes in S also removes all the uncut edges, the remaining graph is bipartite. Moreover, $y(E[S]) = \alpha$.

By the lightness assumption and the weight space normalization, the optimal solution α^* satisfies $\alpha^* \leq y(E)/\Omega(\log(n)) = 1/\Omega(\log(n))$. Running the $O(\sqrt{\log(n)})$ approximation algorithm then outputs a solution with value $\alpha \leq O(\sqrt{\log(n)}) \alpha^* \leq 1/\Omega(\sqrt{\log(n)}) = o(1)$. This therefore leads to an approximation guarantee of

$$R(w) = \left(1 + \frac{1}{\rho}\right) (1 - \alpha) + 2\alpha = 1 + \frac{1}{\rho} + \alpha \left(1 - \frac{1}{\rho}\right) \leq 1 + \frac{1}{\rho} + o(1). \quad \blacktriangleleft$$

References

- 1 Amit Agarwal, Moses Charikar, Konstantin Makarychev, and Yury Makarychev. $o(\sqrt{\log(n)})$ approximation algorithms for min uncut, min 2cnf deletion, and directed cut problems. In *Symposium on the Theory of Computing*, 2005.
- 2 Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, and Bertrand Simon. Online metric algorithms with untrusted predictions. In *International Conference on Machine Learning*, pages 345–355. PMLR, 2020.
- 3 Antonios Antoniadis, Themis Gouleakis, Pieter Kleer, and Pavel Kolev. Secretary and online matching problems with machine learned advice. *Advances in Neural Information Processing Systems*, 33:7933–7944, 2020.
- 4 Sanjeev Arora, Béla Bollobás, and László Lovász. Proving integrality gaps without knowing the linear program. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 313–322. IEEE, 2002.
- 5 Stephan Artmann, Robert Weismantel, and Rico Zenklusen. A strongly polynomial algorithm for bimodular integer linear programming. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1206–1219, 2017.
- 6 R Balasubramanian, Michael R Fellows, and Venkatesh Raman. An improved fixed-parameter algorithm for vertex cover. *Information Processing Letters*, 65(3):163–168, 1998.
- 7 Étienne Bamas, Andreas Maggiori, and Ola Svensson. The primal-dual method for learning augmented algorithms. *Advances in Neural Information Processing Systems*, 33:20083–20094, 2020.
- 8 Reuven Bar-Yehuda and Shimon Even. A linear-time approximation algorithm for the weighted vertex cover problem. *Journal of Algorithms*, 2(2):198–203, 1981.
- 9 Reuven Bar-Yehuda and Shimon Even. A local-ratio theorem for approximating the weighted vertex cover problem. Technical report, Computer Science Department, Technion, 1983.
- 10 Abbas Bazzi, Samuel Fiorini, Sebastian Pokutta, and Ola Svensson. No small linear program approximates vertex cover within a factor $2 - \varepsilon$. *Mathematics of Operations Research*, 44(1):147–172, 2019.
- 11 Adrian Bock, Yuri Faenza, Carsten Moldenhauer, and Andres Jacinto Ruiz-Vargas. Solving the stable set problem in terms of the odd cycle packing number. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science (FSTTCS 2014)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.
- 12 Jonathan F Buss and Judy Goldsmith. Nondeterminism within p. *SIAM Journal on Computing*, 22(3):560–572, 1993.
- 13 Jianer Chen, Iyad A Kanj, and Weijia Jia. Vertex cover: further observations and further improvements. *Journal of Algorithms*, 41(2):280–301, 2001.
- 14 Jianer Chen, Lihua Liu, and Weijia Jia. Improvement on vertex cover for low-degree graphs. *Networks: An International Journal*, 35(4):253–259, 2000.
- 15 Michele Conforti, Samuel Fiorini, Tony Huynh, Gwenaél Joret, and Stefan Weltge. The stable set problem in graphs with bounded genus and bounded odd cycle packing number. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2896–2915. SIAM, 2020.
- 16 Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of mathematics*, pages 439–485, 2005.
- 17 Rodney G Downey and Michael R Fellows. Fixed-parameter tractability and completeness. In *Complexity Theory: Current Research*, pages 191–225, 1992.
- 18 Rodney G Downey and Michael Ralph Fellows. *Parameterized complexity*. Springer Science & Business Media, 2012.
- 19 Samuel Fiorini, Gwenaél Joret, Stefan Weltge, and Yelena Yuditsky. Integer programs with bounded subdeterminants and two nonzeros per row. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 13–24. IEEE, 2022.

- 20 Wayne Goddard and Honghai Xu. Fractional, circular, and defective coloring of series-parallel graphs. *Journal of Graph Theory*, 81(2):146–153, 2016.
- 21 Ervin Györi, Alexandr V Kostochka, and Tomasz Łuczak. Graphs without short odd cycles are nearly bipartite. *Discrete Mathematics*, 163(1-3):279–284, 1997.
- 22 Eran Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. *SIAM Journal on Computing*, 31(5):1608–1623, 2002.
- 23 Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- 24 Dorit S Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on computing*, 11(3):555–556, 1982.
- 25 Dorit S Hochbaum. Efficient bounds for the stable set, vertex cover and set packing problems. *Discrete Applied Mathematics*, 6(3):243–254, 1983.
- 26 George Karakostas. A better approximation ratio for the vertex cover problem. In *International Colloquium on Automata, Languages, and Programming*, pages 1043–1050. Springer, 2005.
- 27 Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- 28 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2 - \epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.
- 29 Stefan Kratsch and Magnus Wahlström. Compression via matroids: a randomized polynomial kernel for odd cycle transversal. *ACM Transactions on Algorithms (TALG)*, 10(4):1–15, 2014.
- 30 Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- 31 Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Online scheduling via learned weights. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1859–1877. SIAM, 2020.
- 32 Lap Chi Lau, Ramamoorthi Ravi, and Mohit Singh. *Iterative methods in combinatorial optimization*, volume 46. Cambridge University Press, 2011.
- 33 Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. *Journal of the ACM (JACM)*, 68(4):1–25, 2021.
- 34 Burkhard Monien and Ewald Speckenmeyer. Ramsey numbers and an approximation algorithm for the vertex cover problem. *Acta Informatica*, 22(1):115–123, 1985.
- 35 George L Nemhauser and Leslie Earl Trotter. Vertex packings: structural properties and algorithms. *Mathematical Programming*, 8(1):232–248, 1975.
- 36 Rolf Niedermeier and Peter Rossmanith. Upper bounds for vertex cover further improved. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 561–570. Springer, 1999.
- 37 Rolf Niedermeier and Peter Rossmanith. On efficient fixed-parameter algorithms for weighted vertex cover. *Journal of Algorithms*, 47(2):63–77, 2003.
- 38 Christos Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 229–234, 1988.
- 39 Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ml predictions. *Advances in Neural Information Processing Systems*, 31, 2018.
- 40 Bruce Reed, Kaleigh Smith, and Adrian Vetta. Finding odd cycle transversals. *Operations Research Letters*, 32(4):299–301, 2004.
- 41 Tim Roughgarden. *Beyond the worst-case analysis of algorithms*. Cambridge University Press, 2021.
- 42 Edward R Scheinerman and Daniel H Ullman. *Fractional graph theory: a rational approach to the theory of graphs*. Courier Corporation, 2011.
- 43 Alexander Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, 1998.
- 44 Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.

20:16 Round and Bipartize for Vertex Cover Approximation

- 45 Mohit Singh. Integrality gap of the vertex cover linear programming relaxation. *Operations Research Letters*, 47(4):288–290, 2019.
- 46 Ulrike Stege and Michael Ralph Fellows. An improved fixed parameter tractable algorithm for vertex cover. *Technical report/Departement Informatik, ETH Zürich*, 318, 1999.

A Integrality Gap and Fractional Chromatic Number

We focus in this section on proving tight bounds for the integrality gap of 3-colorable graphs. A key result that we use in this section is given by Singh in [45], which relates the integrality gap with the fractional chromatic number of a graph. The latter is denoted as $\chi^f(\mathcal{G})$ and is defined as the optimal solution of the following primal-dual linear programming pair. We denote by $\mathcal{I} \subset 2^V$ the set of all independent sets of the graph \mathcal{G} . Solving these linear programs is however NP-hard because of the possible exponential number of independent sets.

$$\begin{array}{ll} \min \sum_{I \in \mathcal{I}} y_I & \max \sum_{v \in V} z_v \\ \sum_{I \in \mathcal{I}, v \in I} y_I \geq 1 \quad \forall v \in V & \sum_{v \in I} z_v \leq 1 \quad \forall I \in \mathcal{I} \\ y_I \geq 0 \quad \forall I \in \mathcal{I} & z_v \geq 0 \quad \forall v \in V \end{array}$$

Note that $\chi^f(\mathcal{G}) = 2$ if and only if \mathcal{G} is bipartite.

► **Theorem 18** (Singh, [45]). *Let $\mathcal{G} = (V, E)$ be a graph. The integrality gap of the vertex cover linear programming relaxation $P(\mathcal{G})$ is:*

$$IG(\mathcal{G}) = 2 - \frac{2}{\chi^f(\mathcal{G})},$$

where $\chi^f(\mathcal{G})$ is the fractional chromatic number of the graph \mathcal{G} .

We first focus on graphs with the existence of a single vertex whose removal produces a bipartite graph. The following theorem generalizes the result given for the cycle graph in [4, 42] and turns out to be the same formula as for series-parallel graphs [20].

► **Theorem 19.** *Let $\mathcal{G} = (V, E)$ be a non-bipartite graph and $v_p \in V$ such that $\mathcal{G} \setminus v_p = (A \cup B, E')$ is bipartite. Then,*


$$\chi^f(\mathcal{G}) = 2 + \frac{1}{\rho - 1},$$

where $2\rho - 1$ is the odd girth of \mathcal{G} .

Proof of Theorem 19. We prove this theorem by constructing feasible primal and dual solutions of objective value $2 + 1/(\rho - 1)$. By strong duality, these two solutions are then optimal for their respective linear programs, hence proving the theorem.

We first construct the dual solution. Let \mathcal{C} be the set of all the shortest odd cycles of \mathcal{G} . For any such cycle $C \in \mathcal{C}$, define the dual solution $z^C \in \mathbb{R}^V$ by

$$z_v^C = \begin{cases} 1/(\rho - 1) & \text{if } v \in C \\ 0 & \text{if } v \in V \setminus C \end{cases}$$



■ **Figure 5** An optimal dual solution constructed in the proof of Theorem 19. Each node on a shortest odd cycle is assigned a fractional value of $1/(\rho - 1)$.

This solution is feasible since any independent set in an odd cycle of length $2\rho - 1$ has size at most $\rho - 1$. Indeed, fix an independent set $I \in \mathcal{I}$, then:

$$\sum_{v \in I} z_v^C = \sum_{v \in I \cap C} \frac{1}{\rho - 1} = \frac{|I \cap C|}{\rho - 1} \leq 1.$$

Moreover, the objective value of this solution is:

$$\sum_{v \in V} z_v^C = \sum_{v \in C} \frac{1}{\rho - 1} = \frac{2\rho - 1}{\rho - 1} = 2 + \frac{1}{\rho - 1}.$$

Let us now construct the primal solution. We will do so by constructing $2\rho - 1$ independent sets $I_k \in \mathcal{I}$ and assigning to each of them a fractional value of $y(I_k) = 1/(\rho - 1)$. All the other independent sets are assigned value zero. We split the bipartite graph $\mathcal{G} \setminus v_p$ into the layers

$$\mathcal{L}_i := \{v \in A \cup B \mid d(N_A(v_p), v) = i\} \quad \text{for } i \in \{0, \dots, l\}.$$


as explained in (2). As a reminder, any shortest odd cycle corresponds to a path between $\mathcal{L}_0 = N_A(v_p)$ and $\mathcal{L}_{2\rho-3} \cap N_B(v_p)$. The original vertex set V is thus decomposed into $\{v_p\} \cup \mathcal{L}_0 \cup \dots \cup \mathcal{L}_l$, where each layer is an independent set and only has edges going out to v_p or its two neighbouring layers.

Let us first focus on the subgraph consisting of the vertices in $\{v_p\} \cup \bigcup_{i=0}^{2\rho-3} \mathcal{L}_i$, where any shortest odd cycle has exactly one vertex per layer (per abuse of notation, we say that $\{v_p\}$ is also a layer in this situation). For convenience of indexing, we rename these layers as $\tilde{\mathcal{L}}_1, \dots, \tilde{\mathcal{L}}_{2\rho-1}$ where $\tilde{\mathcal{L}}_1 = v_p$ and $\tilde{\mathcal{L}}_i = \mathcal{L}_{i-2}$ for $i > 1$. We now create $2\rho - 1$ independent sets on this subgraph in the following way. The first independent set is defined as $U_1 = \tilde{\mathcal{L}}_1 \cup \tilde{\mathcal{L}}_4 \cup \tilde{\mathcal{L}}_6 \dots \cup \tilde{\mathcal{L}}_{2\rho-2}$, where we take the first layer $\tilde{\mathcal{L}}_1$, skip two before taking the next one and then continue by taking the remaining layers alternatingly (hence always skipping one), see Figure 6. Note that the layer following $\tilde{\mathcal{L}}_{2\rho-1}$ is assumed to be $\tilde{\mathcal{L}}_1$. This procedure generates in fact a distinct independent set by starting at $\tilde{\mathcal{L}}_k$ for any $k \in [2\rho - 1]$ and we denote the corresponding independent set by U_k . Notice that each layer is contained in exactly $\rho - 1$ of the constructed independent sets $\{U_k \mid k \in [2\rho - 1]\}$.

We now focus on the subgraph consisting of the vertices in $\bigcup_{i>2\rho-3} \mathcal{L}_i$. We can construct two different independent sets there by taking either the odd or even indexed layers, i.e.

$$R_1 := \bigcup_{i \text{ odd}, i>2\rho-3} \mathcal{L}_i \quad \text{and} \quad R_2 := \bigcup_{i \text{ even}, i>2\rho-3} \mathcal{L}_i.$$

20:18 Round and Bipartize for Vertex Cover Approximation



■ **Figure 6** The $2\rho - 1$ independent sets I_k constructed in the optimal primal solution. The blue nodes correspond to $\{U_k \mid k \in [2\rho - 1]\}$, whereas the orange nodes correspond to $\{R_1, R_2\}$.

We now define our final $2\rho - 1$ independent sets on the full graph as:

$$I_k := \begin{cases} U_k \cup R_1 & \text{if } v_p \notin U_k \\ U_k \cup R_2 & \text{if } v_p \in U_k \end{cases} \quad \forall k \in [2\rho - 1].$$

These are in fact independent sets: in the first case, the first layer in R_1 is $\mathcal{L}_{2\rho-1}$ whereas the last layer in U_k has index at most $2\rho - 3$, meaning that there are no two neighbouring layers. In the second case, since $v_p \in U_k$, we have that $\mathcal{L}_{2\rho-3} \notin U_k$, by construction of U_k . The last layer in U_k thus has index at most $2\rho - 4$, whereas the first layer in R_2 is $\mathcal{L}_{2\rho-2}$, meaning again that there are no two neighbouring layers. In addition, there is no edge between v_p and R_2 , because the only even indexed layer having edges sent to v_p is $\mathcal{L}_0 = N_A(v_p)$.

We now define our primal solution as

$$y(I_k) = \frac{1}{\rho - 1} \quad \forall k \in [2\rho - 1],$$

and $y(I) = 0$ for every other independent set $I \in \mathcal{I}$. We now show this is a feasible solution, i.e. that every vertex $v \in V$ belongs to at least $\rho - 1$ independent sets in $\{I_k \mid k \in [2\rho - 1]\}$. For $v \in \{v_p\} \cup \bigcup_{i=0}^{2\rho-3} \mathcal{L}_i$, such a vertex lies by construction in exactly $\rho - 1$ independent sets $\{U_k \mid k \in [2\rho - 1]\}$, and thus also of $\{I_k \mid k \in [2\rho - 1]\}$. For $v \in \bigcup_{i>2\rho-3} \mathcal{L}_i$, if v belongs to an even indexed layer, then it is contained in $\rho - 1$ of the desired independent sets. If it belongs to an odd indexed layer, then it is contained in ρ of them. Therefore,

$$\sum_{I \in \mathcal{I}, v \in I} y(I) = \sum_{k=1}^{2\rho-1} y(I_k) \mathbf{1}_{\{v \in I_k\}} = \frac{1}{\rho - 1} \sum_{k=1}^{2\rho-1} \mathbf{1}_{\{v \in I_k\}} \geq 1.$$

The objective value of this primal solution is clearly $2 + 1/(\rho - 1)$. We have constructed feasible primal and dual solutions with the same objective value. By strong duality, this finishes the proof of the theorem. ◀

We now consider the case where $\mathcal{G} = (V, E)$ is a graph with chromatic number $\chi(\mathcal{G}) = 3$.

► **Theorem 20.** *Let $\mathcal{G} = (V, E)$ be a 3-colorable graph with color classes $V = V_1 \cup V_2 \cup V_3$. Then,*

$$\chi^f(\mathcal{G}) \leq 2 + \min_{i \in \{1,2,3\}} \frac{1}{\rho_i - 1}$$

where $2\rho_i - 1$ is the odd girth of the contracted graph \mathcal{G}/V_i for each $i \in \{1, 2, 3\}$. Moreover, equality holds if one color class only contains one vertex.

Proof. We prove this theorem by constructing three feasible solutions of value $2 + 1/(\rho_i - 1)$ for each $i \in \{1, 2, 3\}$ to the primal linear program of the fractional chromatic number on the graph \mathcal{G} .

Fix an $i \in \{1, 2, 3\}$ and consider the graph $\tilde{\mathcal{G}} := \mathcal{G}/V_i = (\tilde{V}, \tilde{E})$ with odd girth $2\rho_i - 1$. We denote the contracted node by $\tilde{v} \in \tilde{V}$. Since this graph is bipartite if we were to remove \tilde{v} , we know that its fractional chromatic number is equal to $2 + 1/(\rho_i - 1)$ by Theorem 19. Let $\{\tilde{I}_k, k \in [2\rho_i - 1]\}$ be the independent sets in the support of the optimal primal solution of the graph $\tilde{\mathcal{G}}$ constructed in the proof of this theorem. For each of these independent sets, we extend them to the original graph in the following way:

$$I_k = \begin{cases} \tilde{I}_k & \text{if } \tilde{v} \notin \tilde{I}_k \\ (\tilde{I}_k \setminus \tilde{v}) \cup V_i & \text{if } \tilde{v} \in \tilde{I}_k. \end{cases}$$

In words, if \tilde{v} happens to belong to \tilde{I}_k , we replace it by V_i to get a valid independent set in the original graph. Assigning fractional value $y(I_k) = 1/(\rho_i - 1)$ for every $k \in [2\rho_i - 1]$ yields a feasible primal solution with objective value $2 + 1/(\rho_i - 1)$. Since we can do this for every $i \in \{1, 2, 3\}$, and the optimal minimum value of the primal linear program is at most the objective value of any of these feasible solutions, the proof is finished.

Moreover, this upper bound is in fact tight, since it holds with equality when one of the color classes only contains one vertex by Theorem 19. ◀

B Arbitrary Set to Bipartite: Omitted Proofs

► **Theorem 21.** *Let $\mathcal{G} = (V, E)$ be a graph and $S \subset V$ such that $\mathcal{G} \setminus S = (A \cup B, E')$ is bipartite. For any $w \in Q^W$, the approximation ratio satisfies*

$$R(w) \leq 1 + \alpha \quad \text{with } \alpha \in [0, 1]$$

if the contracted graph \mathcal{G}/S is bipartite.

Proof. The only change with respect to the previous proof is the bound on $w(OPT(\mathcal{G} \setminus S))$ in (5). We denote the contracted graph by $\tilde{\mathcal{G}} = \mathcal{G}/S$ and by v^S the contracted vertex. Suppose $\tilde{\mathcal{G}}$ admits the bipartition $(\tilde{A} \cup \tilde{B}, \tilde{E})$ and assume without loss of generality that $v^S \in \tilde{A}$. Note that $\tilde{E} = E' \cup \delta(v^S)$.

Any feasible cover of $\mathcal{G} \setminus S$ needs to count the dual value of every edge in E' at least once. Taking the cover $\tilde{A} \setminus v^S$ counts every edge in E' exactly once, showing that $w(OPT(\mathcal{G} \setminus S)) = y(E')$. Hence, using $y(E) = \alpha + y(\delta(S)) + y(E') = 1$, we get

$$R(w) \leq 2\alpha + y(\delta(S)) + y(E') = 1 + \alpha. \quad \blacktriangleleft$$

20:20 Round and Bipartize for Vertex Cover Approximation

► **Theorem 22.** *Let $\alpha \in [0, 1]$ and let $\rho \geq 2$. There exists a non-bipartite graph $\mathcal{G} = (V, E)$, with weights $(y, w) \in Q^{Y, W}$, and a set $S \subset V$ with $y(E[S]) = \alpha$, where \mathcal{G}/S has odd girth $2\rho - 1$ and which satisfies*

$$R(w) = \left(1 + \frac{1}{\rho}\right) (1 - \alpha) + 2\alpha.$$

Proof. An example of such a graph can be constructed as follows. We first construct \mathcal{G}/S : take an odd cycle of length $2\rho - 1$ with a distinguished node v^S and assign dual value $(1 - \alpha)/\rho$ to both edges incident to it. Alternatively assign dual values 0 and $(1 - \alpha)/\rho$ along the odd cycle for the remaining edges. In order to construct \mathcal{G} , replace v^S by a triangle S with dual edges set to $\alpha, 0$ and 0 , where the two previous incident edges to v^S are adjacent to the endpoints of the edge with value α . Note that we replace it with a triangle instead of a single edge in order to avoid \mathcal{G} becoming bipartite. Similarly to the proof of Theorem 12, one can check that

$$w(S) = 2\alpha + \frac{2(1 - \alpha)}{\rho}; \quad w(OPT(\mathcal{G} \setminus S)) = \frac{(1 - \alpha)(\rho - 1)}{\rho}; \quad w(OPT(\mathcal{G})) = 1.$$

Therefore,

$$R(w) = 2\alpha + \frac{2(1 - \alpha)}{\rho} + \frac{(1 - \alpha)(\rho - 1)}{\rho} = \left(1 + \frac{1}{\rho}\right) (1 - \alpha) + 2\alpha. \quad \blacktriangleleft$$

► **Theorem 23.** *Let $\alpha \in [0, 1]$. There exists a non-bipartite graph $\mathcal{G} = (V, E)$, with weights $(y, w) \in Q^{Y, W}$, and a set $S \subset V$ with $y(E[S]) = \alpha$, where \mathcal{G}/S is bipartite and which satisfies*

$$R(w) = 1 + \alpha.$$

Proof. Let \mathcal{G} be an arbitrary odd cycle. Consider an arbitrary edge $(u, v) \in E$ and assign it dual value α . The set S is defined to be $S = \{u, v\}$. Assign dual value zero to the edge $(u, w) \in E$, where w is the second neighbour of u in the cycle. For the remaining edges, arbitrarily assign dual values, while ensuring that they sum up to $1 - \alpha$. The fact that one edge is equal to zero is necessary in order to get the exact formula $w(OPT(\mathcal{G})) = 1$, a feasible cover showing $w(OPT(\mathcal{G})) \leq 1$ being the following: take both endpoints of the edge (u, w) and take remaining vertices alternatively (hence always skipping one) along the odd cycle. All the edges are counted once, except for (u, w) , which is counted twice but has value zero. Moreover, $w(S) = 2\alpha + y(\delta(S))$ and $w(OPT(\mathcal{G} \setminus S)) = y(E')$, where E' is the edge set of the bipartite graph $\mathcal{G} \setminus S$. Therefore,

$$R(w) = 2\alpha + y(\delta(S)) + y(E') = 1 + \alpha. \quad \blacktriangleleft$$