

A $2\ell k$ Kernel for ℓ -Component Order Connectivity*

Mithilesh Kumar¹ and Daniel Lokshtanov²

- 1 Department of Informatics, University of Bergen, Norway
Mithilesh.Kumar@ii.uib.no
- 2 Department of Informatics, University of Bergen, Norway
daniello@ii.uib.no

Abstract

In the ℓ -COMPONENT ORDER CONNECTIVITY problem ($\ell \in \mathbb{N}$), we are given a graph G on n vertices, m edges and a non-negative integer k and asks whether there exists a set of vertices $S \subseteq V(G)$ such that $|S| \leq k$ and the size of the largest connected component in $G - S$ is at most ℓ . In this paper, we give a kernel for ℓ -COMPONENT ORDER CONNECTIVITY with at most $2\ell k$ vertices that takes $n^{\mathcal{O}(\ell)}$ time for every constant ℓ . On the way to obtaining our kernel, we prove a generalization of the q -Expansion Lemma to weighted graphs. This generalization may be of independent interest.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Parameterized algorithms, Kernel, Component Order Connectivity, Max-min allocation, Weighted expansion

Digital Object Identifier 10.4230/LIPIcs.IPEC.2016.20

1 Introduction

In the classic VERTEX COVER problem, the input is a graph G and integer k , and the task is to determine whether there exists a vertex set S of size at most k such that every edge in G has at least one endpoint in S . Such a set is called a *vertex cover* of the input graph G . An equivalent definition of a vertex cover is that every connected component of $G - S$ has at most 1 vertex. This view of the VERTEX COVER problem gives rise to a natural generalization: can we delete at most k vertices from G such that every connected component in the resulting graph has at most ℓ vertices? Here we study this generalization. Formally, for every integer $\ell \geq 1$, we consider the following problem, called ℓ -COMPONENT ORDER CONNECTIVITY (ℓ -COC).

ℓ -COMPONENT ORDER CONNECTIVITY (ℓ -COC)

Input: A graph G on n vertices and m edges, and a positive integer k .

Task: determine whether there exists a set $S \subseteq V(G)$ such that $|S| \leq k$ and the maximum size of a component in $G - S$ is at most ℓ .

The set S is called an ℓ -COC *solution*. For $\ell = 1$, ℓ -COC is just the VERTEX COVER problem. Aside from being a natural generalization of VERTEX COVER, the family $\{\ell\text{-COC} : \ell \geq 1\}$ of problems can be thought of as a vulnerability measure of the graph G - how many vertices

* The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. 306992 and the Beating Hardness by Pre-processing grant funded by the Bergen Research Foundation.



of G have to fail for the graph to break into small connected components? For a study of ℓ -COC from this perspective see the survey of Gross et al. [12].

From the work of Lewis and Yannakakis [16] it immediately follows that ℓ -COC is NP-complete for every $\ell \geq 1$. This motivates the study of ℓ -COC within paradigms for coping with NP-hardness, such as approximation algorithms [22], exact exponential time algorithms [11], parameterized algorithms [5, 7] and kernelization [14, 17]. The ℓ -COC problems have (for some values of ℓ) been studied within all four paradigms, see the related work section.

In this work we focus on ℓ -COC from the perspective of parameterized complexity and kernelization. Our main result is an algorithm that given an instance (G, k) of ℓ -COC, runs in polynomial time, and outputs an equivalent instance (G', k') such that $k' \leq k$ and $|V(G')| \leq 2\ell k$. This is called a *kernel* for ℓ -COC with $2\ell k$ vertices. Our kernel significantly improves over the previously best known kernel with $O(\ell k(k + \ell))$ vertices by Drange et al. [8]. Indeed, for $\ell = 1$ our kernel matches the size of the smallest known kernel for VERTEX COVER [3] that is based on the classic theorem of Nemhauser and Trotter [18].

Related Work. 1-COC, better known as VERTEX COVER, is extremely well studied from the perspective of approximation algorithms [22, 6], exact exponential time algorithms [10, 19, 24], parameterized algorithms [5, 4] and kernelization [3, 18]. The kernel with $2k$ vertices for VERTEX COVER is considered one of the fundamental results in the field of kernelization. The 2-COC problem is also well studied, and has been considered under several different names. The problem, or rather the dual problem of finding a largest possible set S that induces a subgraph in which every connected component has order at most 2, was first defined by Yannakakis [25] under the name Dissociation Set. The problem has attracted attention in exact exponential time algorithms [13, 23], the fastest currently known algorithm [23] has running time $O(1.3659^n)$. 2-COC has also been studied from the perspective of parameterized algorithms [2, 20] (under the name VERTEX COVER P_3) as well as approximation algorithms [21]. The fastest known parameterized algorithm, due to Chang et al. [2] has running time $1.7485^k n^{O(1)}$, while the best approximation algorithm, due to Tu and Zhou [21] has factor 2.

For the general case of ℓ -COC, $\ell \geq 1$, Drange et al. [8] gave a simple parameterized algorithm with running time $(\ell + 1)^k n^{O(1)}$, and a kernel with $O(k\ell(\ell + k))$ vertices. The parameterized algorithm of Drange et al. [8] can be improved to $(\ell + 0.0755)^k n^{O(1)}$ by reducing to the $(\ell + 1)$ -HITTING SET problem, and applying the iterative compression based algorithm for $(\ell + 1)$ -HITTING SET due to Fomin et al. [9]. The reduction to $(\ell + 1)$ -HITTING SET, coupled with the simple factor $(\ell + 1)$ -approximation algorithm for $(\ell + 1)$ -HITTING SET [22] immediately also yields an $(\ell + 1)$ -approximation algorithm for ℓ -COC. There has also been some work on ℓ -COC when the input graph is restricted to belong to a graph class, for a discussion of this work see [8].

Comparing the existing results with our work, we see that our kernel improves over the kernel of Drange et al. [8] from at most $O(k\ell(\ell + k))$ vertices to at most $2k\ell$ vertices. Our kernel is also the first kernel with a linear number of vertices for every fixed $\ell \geq 2$.

Our Methods. Our kernel for ℓ -COC hinges on the concept of a *reducible pair* of vertex sets. Essentially (*this is not the formal definition used in the paper!*), a reducible pair is a pair (X, Y) of disjoint subsets of $V(G)$ such that $N(Y) \subseteq X$, every connected component of $G[Y]$ has size at most ℓ , and every solution S to G has to contain at least $|X|$ vertices from $G[X \cup Y]$. If a reducible pair is identified, it is easy to see that one might just as well pick all of X into the solution S , since any solution has to pay $|X|$ inside $G[X \cup Y]$, and after X

is deleted, Y breaks down into components of size at most ℓ and is completely eliminated from the graph.

At this point there are several questions. (a) How does one argue that a reducible pair is in fact reducible? That is, how can we prove that any solution has to contain at least $|X|$ vertices from $X \cup Y$? (b) How big does G have to be compared to k before we can assert the existence of a reducible pair? Finally, (c) even if we can assert that G contains a reducible pair, how can we find one in polynomial time?

To answer (a) we restrict ourselves to reducible pairs with the additional property that each connected component C of $G[Y]$ can be assigned to a vertex $x \in N(C)$, such that for every $x \in X$ the total size of the components assigned to x is at least ℓ . Then x together with the components assigned to it form a set of size at least $\ell + 1$ and have to contain a vertex from the solution. Since we obtain such a connected set for each $x \in X$, the solution has to contain at least $|X|$ vertices from $X \cup Y$. Again we remark that this definition of a reducible pair is local to this section, and not the one we actually end up using.

To answer (b) we first try to use the q -Expansion Lemma (see [5]), a tool that has found many uses in kernelization. Roughly speaking the Expansion Lemma says the following: if $q \geq 1$ is an integer and H is a bipartite graph with bipartition (A, B) and B is at least q times larger than A , then one can find a subset X of A and a subset Y of B such that $N(Y) \subseteq X$, and an assignment of each vertex $y \in Y$ to a neighbor x of y , such that every vertex x in X has at least q vertices in Y assigned to it.

Suppose now that the graph does have an ℓ -COC solution S of size at most k , and that $V(G) \setminus S$ is sufficiently large compared to S . The idea is to apply the Expansion Lemma to the bipartite graph H , where the A side of the bipartition is S and the B side has one vertex for each connected component of $G - S$. We put an edge in H between a vertex v in S and a vertex corresponding to a component C of $G - S$ if there is an edge between v and C in G . If $G - S$ has at least $|S| \cdot \ell$ connected components, we can apply the ℓ -Expansion Lemma on H , and obtain a set $X \subseteq S$, and a collection \mathcal{Y} of connected components of $G - X$ satisfying the following properties. Every component $C \in \mathcal{Y}$ satisfies $N(C) \subseteq X$ and $|C| \leq \ell$. Furthermore, there exists an assignment of each connected component C to a vertex $x \in N(C)$, such that every $x \in X$ has at least ℓ components assigned to it. Since x has at least ℓ components assigned to it, the total size of the components assigned to x is at least ℓ . But then, X and $Y = \bigcup_{C \in \mathcal{Y}} C$ form a reducible pair, giving an answer to question (b). Indeed, this argument can be applied whenever the number of components of $G - S$ is at least $\ell \cdot |S|$. Since each component of $G - S$ has size at most ℓ , this means that the argument can be applied whenever $|V(G) \setminus S| \geq \ell^2 \cdot |S| \geq \ell^2 k$.

Clearly this argument fails to yield a kernel of size $2\ell k$, because it is only applicable when $|V(G)| = \Omega(\ell^2 k)$. At this point we note that the argument above is extremely wasteful in one particular spot: we used the *number* of components assigned to x to lower bound the *total size* of the components assigned to x . To avoid being wasteful, we prove a new variant of the Expansion Lemma, where the vertices on the B side of the bipartite graph H have non-negative integer weights. This new Weighted Expansion lemma states that if $q, W \geq 1$ are integers, H is a bipartite graph with bipartition (A, B) , every vertex in B has a non-negative integer weight which is at most W , and the total weight of B is at least $(q + W - 1) \cdot |A|$, then one can find a subset X of A and a subset Y of B such that $N(Y) \subseteq X$, and an assignment of each vertex $y \in Y$ to a neighbor x of y , such that for every vertex in X , the total weight of the vertices assigned to it is at least q . The proof of the Weighted Expansion Lemma is based on a combination of the usual, unweighted Expansion Lemma with a variant of an argument by Bezáková and Dani [1] to round the linear program for

Max-min Allocation of goods to customers.

Having the Weighted Expansion Lemma at hand we can now repeat the argument above for proving the existence of a reducible pair, but this time, when we build H , we can give the vertex corresponding to a component C of $G - S$ weight $|C|$, and apply the Weighted Expansion Lemma with $q = \ell$ and $W = \ell$. Going through the argument again, it is easy to verify that this time the existence of a reducible pair is guaranteed whenever $|V(G) \setminus S| \leq (2\ell - 1)k$, that is when $|V(G)| \geq 2\ell k$.

We are now left with question (c) - the issue of how to *find* a reducible pair in polynomial time. Indeed, the proof of existence crucially relies on the knowledge of an (optimal) solution S . To find a reducible pair we use the linear programming relaxation of the ℓ -COC problem. We prove that an optimal solution to the LP-relaxation has to highlight every reducible pair (X, Y) , essentially by always setting all the variables corresponding to X to 1 and the variables corresponding to Y to 0. For VERTEX COVER (i.e 1-COC), the classic Nemhauser Trotter Theorem [18] implies that we may simply include all the vertices whose LP variable is set to 1 into the solution S . For ℓ -COC with $\ell \geq 2$ we are unable to prove the corresponding statement. We are however, able to prove that if a reducible pair (X, Y) exists, then X (essentially) has to be assigned 1 and Y (essentially) has to be assigned 0. We then give a polynomial time algorithm that extracts X and Y from the vertices assigned 1 and 0 respectively by the optimal linear programming solution. Together, the arguments (b) and (c) yield the kernel with $2\ell k$ vertices. We remark that to the best of our knowledge, after the kernel for Vertex Cover [3] our kernel is the first example of a kernelization algorithm based on linear programming relaxations.

Overview of the paper. In Section 2 we recall basic definitions and set up notations. The kernel for ℓ -COC is proved in Sections 3, 4 and 5. In Section 3 we prove the necessary adjustment of the results on Max-Min allocation of Bezáková and Dani [1] that is suitable to our needs. In Section 4 we state and prove our new Weighted Expansion Lemma, and in Section 5 we combine all our results to obtain the kernel.

2 Preliminaries

Let \mathbb{N} denote the set of positive integers $\{0, 1, 2, \dots\}$. For any non-zero $t \in \mathbb{N}$, $[t] := \{1, 2, \dots, t\}$. We denote a constant function $f : X \rightarrow \mathbb{N}$ such that for all $x \in X$, $f(x) = c$, by $f = c$. For any function $f : X \rightarrow \mathbb{N}$ and a constant $c \in \mathbb{N}$, we define the function $f + c : X \rightarrow \mathbb{N}$ such that for all $x \in X$, $(f + c)(x) = f(x) + c$. We use the same symbol f to denote the restriction of f over a subset of its domain, X . For a set $\{v\}$ containing a single element, we simply write v . A vertex $u \in V(G)$ is said to be incident on an edge $e \in E(G)$ if u is one of the endpoints of e . A pair of edges $e, e' \in E(G)$ are said to be adjacent if there is a vertex $u \in V(G)$ such that u is incident on both e and e' . For any vertex $u \in V(G)$, by $N(u)$ we denote the set of neighbors of u i.e. $N(u) := \{v \in V(G) \mid uv \in E(G)\}$. For any subgraph $X \subseteq G$, by $N(X)$ we denote the set of neighbors of vertices in X outside X , i.e. $N(X) := (\bigcup_{u \in X} N(u)) \setminus X$. A pair of vertices $u, v \in V(G)$ are called *twins* if $N(u) = N(v)$. An induced subgraph on $X \subseteq V(G)$ is denoted by $G[X]$.

A *path* P is a graph, denoted by a sequence of vertices $v_1 v_2 \dots v_t$ such that for any $i, j \in [t]$, $v_i v_j \in E(P)$ if and only if $|i - j| = 1$. A *cycle* C is a graph, denoted either by a sequence of vertices $v_1 v_2 \dots v_t$ or by a sequence of edges $e_1 e_2 \dots e_t$, such that for any $i, j \in [t]$ $u_i u_j \in E(C)$ if and only if $|i - j| = 1 \pmod t$ or in terms of edges, for any $i, j \in [t]$, e_i is adjacent to e_j if and only if $|i - j| = 1 \pmod t$. The *length* of a path(cycle) is the number

of edges in the path(cycle). A *triangle* is a cycle of length 3. In G , for any pair of vertices $u, v \in V(G)$ $\text{dist}(u, v)$ represents the length of a shortest path between u and v . A *tree* is a connected graph that does not contain any cycle. A *rooted tree* T is a tree with a special vertex r called the root of T . With respect to r , for any edge $uv \in E(T)$ we say that v is a child of u (equivalently u is parent of v) if $\text{dist}(u, r) < \text{dist}(v, r)$. A *forest* is a collection of trees. A *rooted forest* is a collection of rooted trees. A *clique* is a graph that contains an edge between every pair of vertices. A vertex cover of a graph is a set of vertices whose removal makes the graph edgeless.

Fixed Parameter Tractability. A *parameterized problem* Π is a subset of $\Sigma^* \times \mathbb{N}$. A parameterized problem Π is said to be *fixed parameter tractable* (FPT) if there exists an algorithm that takes as input an instance (I, k) and decides whether $(I, k) \in \Pi$ in time $f(k) \cdot n^c$, where n is the length of the string I , $f(k)$ is a computable function depending only on k and c is a constant independent of n and k .

A *kernel* for a parameterized problem Π is an algorithm that given an instance (T, k) runs in time polynomial in $|T|$, and outputs an instance (T', k') such that $|T'|, k' \leq g(k)$ for a computable function g and $(T, k) \in \Pi$ if and only if $(T', k') \in \Pi$. For a comprehensive introduction to FPT algorithms and kernels, we refer to the book by Cygan et al. [5].

A *data reduction rule*, or simply, reduction rule, for a parameterized problem Q is a function $\phi : \Sigma^* \times \mathbb{N} \rightarrow \Sigma^* \times \mathbb{N}$ that maps an instance (I, k) of Q to an equivalent instance (I', k') of Q such that ϕ is computable in time polynomial in $|I|$ and k . We say that two instances of Q are *equivalent* if $(I, k) \in Q$ if and only if $(I', k') \in Q$; this property of the reduction rule ϕ , that it translates an instance to an equivalent one, is referred as the *safeness* of the reduction rule.

3 Max-min Allocation

We will now view a bipartite graph $G := ((A, B), E)$ as a relationship between “customers” represented by the vertices in A and “items” represented by the vertices in B . If the graph is supplied with two functions $w_a : A \rightarrow \mathbb{N}$ and $w_b : B \rightarrow \mathbb{N}$, we treat these functions as a “demand function” and a “capacity” function, respectively. That is, we consider each item $v \in B$ to have value $w_b(v)$, and every customer $u \in A$ wants to be assigned items worth at least $w_a(u)$. An edge between $u \in A$ and $v \in B$ means that the item v can be given to u .

A weight function $f : E(G) \rightarrow \mathbb{N}$ describes an assignment of items to customers, provided that the items can be “divided” into pieces and the pieces can be distributed to different customers. However this “division” should not create more value than the original value of the items. Formally we say that the weight function *satisfies* the capacity constraint $w_b(v)$ of $v \in B$ if $\sum_{uv \in E(G)} f(uv) \leq w_b(v)$. The weight function satisfies the capacity constraints if it satisfies the capacity constraints of all items $v \in B$.

For each item $u \in A$, we say that f *allocates* $\sum_{uv \in E(G)} f(uv)$ value to u . The weight function f *satisfies* the demand $w_a(u)$ of $u \in A$ if it allocates at least $w_a(u)$ value to u , and f satisfies the demand constraints if it does so for all $u \in A$. In other words, the weight function satisfies the demands if every customer gets items worth at least her demand. The weight function f *over-satisfies* a demand constraint $w_a(u)$ of u if it allocates strictly more than $w_a(u)$ to u .

We will also be concerned with the case where items are indivisible. In particular we say that a weight function $f : E(G) \rightarrow \mathbb{N}$ is *unsplitting* if for every $v \in B$ there is at most one edge $uv \in E(G)$ such that $f(uv) > 0$. The essence of the next few lemmas is that if

we have a (splitting) weight function f of items whose value is at most W , and f satisfies the capacity and demand constraints, then we can obtain in polynomial-time an unsplitting weight function f' that satisfies the capacity constraints and violates the demand constraints by at most $(W - 1)$. In other words we can make a splitting distribution of items unsplitting at the cost of making each customer lose approximately the value of the most expensive item.

Allocating items to customers in such a way as to maximize satisfaction is well studied in the literature. The lemmata 1 and 2 are very similar, both in statement and proof, to the work of Bezáková and Dani [1][Theorem 3.2], who themselves are inspired by Lenstra et al. [15]. However we do not see a way to directly use the results of Bezáková and Dani [1], because we need a slight strengthening of (a special case of) their statement.


► **Lemma 1.** *There exists a polynomial-time algorithm that given a bipartite graph G , a capacity function $w_b : B \rightarrow \mathbb{N}$, a demand function $w_a : A \rightarrow \mathbb{N}$ and a weight function $f : E(G) \rightarrow \mathbb{N}$ that satisfies the capacity and demand constraints, outputs a function $f' : E(G) \rightarrow \mathbb{N}$ such that f' satisfies the capacity and demand constraints and the graph $G_{f'} = (V(G), \{uv \in E(G) \mid f'(uv) > 0\})$ induced on the non-zero weight edges of G is a forest.*

Proof. We start with f and in polynomially many steps, change f into the required function f' . If $G_f = (V(G), \{uv \in E(G) \mid f(uv) > 0\})$ is a forest, then we return $f' = f$. Otherwise, suppose that G_f contains a cycle $C := e_1e_2e_3 \dots e_{2s}$. Proceed as follows. Without loss of generality, suppose $c = f(e_1) = \min\{f(e) \mid e \in C\}$, and note that $c > 0$. Compute the edge weight function $f^* : E \rightarrow \mathbb{R}$ defined as follows. For $e_i \in C$, we define $f^*(e_i) = f(e) - c$ if i is odd, and define $f^*(e_i) = f(e) + c$ if i is even. For $e \notin C$ we define $f^*(e) = f(e)$.

Every vertex of G is incident to either 0 or exactly 2 edges of C . If the vertex v is incident to two edges of C then one of these edges, say e_{2i} , has even index in C , and the other, e_{2i+1} has odd. For the edge e_{2i} we have $f^*(e_{2i}) = f(e_{2i}) + c$ and for e_{2i+1} we have $f^*(e_{2i+1}) = f(e_{2i+1}) - c$. Thus we conclude that for all $v \in V(G)$, $\sum_{u \in N(v)} f^*(uv) = \sum_{u \in N(v)} f(uv)$, and that therefore f^* satisfies the capacity and demand constraints. Furthermore at least one edge that is assigned non-zero weight by f is assigned 0 by f^* and $G_{f^*} = (V(G), \{uv \in E(G) \mid f^*(uv) > 0\})$ has one less cycle than G_f . For a polynomial-time algorithm, repeatedly apply the process described above to reduce the number of edges with non-zero weight, as long as G_{f^*} contains a cycle. ◀

► **Lemma 2.** *There exists a polynomial-time algorithm with the following specifications. It takes as input a bipartite graph $G := ((A, B), E)$, a demand function $w_a : A \rightarrow \mathbb{N}$, a capacity function $w_b : B \rightarrow \mathbb{N}$, an edge weight function $f : E(G) \rightarrow \mathbb{N}$ that satisfies both the capacity and demand constraints, and a vertex $r \in A$. The algorithm outputs an unsplitting edge weight function $h : E(G) \rightarrow \mathbb{N}$ that satisfies the capacity constraints, satisfies the demands $w'_a = w_a - (W - 1)$ where $W = \max_{v \in B} w_b(v)$, and additionally satisfies the demand $w_a(r)$ of r .*

Proof. Without loss of generality the graph $G_f := (V(G), \{uv \in E(G) \mid f(uv) > 0\})$ is a forest. If it is not, we may apply Lemma 1 to f , and obtain a function f' that satisfies the capacity and demand constraints, and such that $G_{f'} = (V(G), \{uv \in E(G) \mid f'(uv) > 0\})$ is a forest. We then rename f' to f . By picking a root in each connected component of G_f we may consider G_f as a rooted forest. We pick the roots as follows, if the component contains the special vertex r , we pick r as root. If the component does not contain r , but contains at least one vertex $u \in A$, we pick that vertex as the root. If the component does not contain any vertices of A then it does not contain any edges and is therefore a single vertex in B , we



■ **Figure 1** Proof of Lemma 1 and 2. Cyclically shift smallest weight in a non-zero weight cycle to obtain a forest. Root each tree in the forest at a vertex in A such that each vertex in B has a parent in A . Assign the value of $v \in B$ to its parent $u \in A$. In this new assignment, a non-root vertex $u \in A$ loses its parent $v_0 \in B$ and $f(v_0u) \leq W - 1$ which explains the cost of making a splitting assignment into an unsplitting assignment.

pick that vertex as root. Thus, every item $v \in B$ that is incident to at least one edge in G_f has a unique parent $u \in A$ in the forest G_f . We define the new weight function h . For every edge $uv \in E(G)$ with $u \in A$ and $v \in B$ we define $h(uv)$ as follows.

$$h(uv) = w_b(v) \text{ if } u \text{ is the parent of } v \text{ in } G_f, \text{ and } h(uv) = 0 \text{ otherwise.}$$

Clearly h is unsplitting and satisfies the capacity constraints. We now prove that h also satisfies the demand constraints w'_a and satisfies the demand constraint $w_a(r)$ of r . Consider the demand constraint $w'_a(u)$ for an arbitrary customer $u \in A$. There are two cases, either u is the root of the component of G_f or it is not. If u is the root, then for every edge $uv \in E(G)$ such that $f(uv) > 0$ we have that $uv \in E(G_f)$ and consequently that u is the parent of v . Hence $h(uv) = w_b(v) \geq f(uv)$, and therefore h satisfies the demand $w_a(u)$ of u . Since $w_a(u) \geq w'_a(u)$, we have that h satisfies the demand $w'_a(u)$. Furthermore, since r is the root of its component this also proves that h satisfies the demand $w_a(r)$.

Consider now the case that u is not the root of its component in G_f . Then u has a unique parent in G_f , call this vertex $v^* \in B$. We first prove that $f(uv^*) \leq w_b(v^*) - 1$. Indeed, since v^* is incident to the edge uv^* we have that v^* has a parent u^* in G_f , and that $u^* \neq u$ because v^* is the parent of u . We have that $f(u^*v^*) + f(uv^*) \leq w_b(v^*)$ and that $f(u^*v^*) \geq 1$, because u^*v^* is an edge in G_f . It follows that $f(uv^*) \leq w_b(v^*) - 1$. We now proceed to proving that h satisfies the demand $w'_a(u)$.

For every edge $uv \in E(G) \setminus \{uv^*\}$ such that $uv \in E(G)$ such that $f(uv) > 0$ we have that $uv \in E(G_f)$ and consequently that u is the parent of v . Hence we have that $h(uv) = w_b(v) \geq f(uv)$. Furthermore $h(uv^*) = 0$ while $f(uv^*) \leq w_b(v^*) - 1 \leq W - 1$. Therefore h satisfies the demand $w'_a(u)$. ◀

4 The Weighted Expansion Lemma

Our kernelization algorithm will use “ q -expansions” in bipartite graphs, a well known tool in kernelization [5]. We begin by stating the definition of a q -expansion and review the facts about them that we will use.

► **Definition 3** (q -expansion). Let $G := ((A, B), E)$ be a bipartite graph. We say that A has q -expansion into B if there is a family of sets $\{V_a \mid V_a \subseteq N(a), |V_a| \geq q, a \in A\}$ such that for any pair of vertices $a_i, a_j \in A, i \neq j, V_{a_i} \cap V_{a_j} = \emptyset$.

► **Definition 4** (Twin graph). For a bipartite graph $G := ((A, B), E)$ with a weight function $w_b : B \rightarrow \mathbb{N}$, the twin graph $T_{AB} := (A, B')$ of G is obtained as follows: B' contains $|w_b(v)|$ twins of every vertex $v \in B$ i.e. $B' := \{v_1, v_2, \dots, v_{w_b(v)} \mid v \in B\}$ and edges in T_{AB} such that for all $v \in B$ and $i \in [w_b(v)], N(v_i) = N(v)$ i.e. $E(T_{AB}) := \{av_i \mid a \in A, v_i \in B', v \in B, av \in E(G)\}$.

► **Lemma 5** ([5]). *Let G be a bipartite graph with bipartition (A, B) . Then there is a q -expansion from A into B if and only if $|N(X)| \geq q|X|$ for every $X \subseteq A$. Furthermore, if there is no q -expansion from A into B , then a set $X \subseteq A$ with $|N(X)| < q|X|$ can be found in polynomial-time.*

► **Lemma 6** (Expansion Lemma [5]). *Let $q \geq 1$ be a positive integer and G be a bipartite graph with vertex bipartition (A, B) such that $|B| \geq q|A|$, and there are no isolated vertices in B . Then there exist nonempty vertex sets $X \subseteq A$ and $Y \subseteq B$ such that there is a q -expansion of X into Y , and no vertex in Y has a neighbor outside X , i.e. $N(Y) \subseteq X$. Furthermore, the sets X and Y can be found in time polynomial in the size of G .*

► **Lemma 7** (Folklore). *There exists a polynomial-time algorithm that given a bipartite graph $G := ((A, B), E)$ and an integer q decides (and outputs in case yes) if there exist sets $X \subseteq A, Y \subseteq B$ such that there is a q -expansion of X into Y .*

Proof. We describe a recursive algorithm. If $A = \emptyset$ or $B = \emptyset$, then output NO and terminate. Otherwise, construct the twin graph T_{BA} with weight function $w : A \rightarrow \mathbb{N}$ where for all $u \in A, w(u) = q$ and let M be a maximum matching in T_{BA} . Consider the graph $G' := (A, B)$ with edge set $E(G') := \{uv, u \in A, v \in B \mid u_i v \in M\}$. Let $A' \subseteq A$ such that for all $u \in A', d_{G'}(u) \geq q$ and let $B' \subseteq B$ such that $B' := \bigcup_{u \in A'} N_{G'}(u)$. If $N(B') \subseteq A'$, then return (A', B') and terminate. Otherwise, recurse on $G[A' \cup (B \setminus N_G(A \setminus A'))]$.

If there are no sets X, Y such that there is a q -expansion of X into Y , then for any pair of sets $A' \subseteq A, B' \subseteq B$ either $N(B') \setminus A' \neq \emptyset$ or $|B'| < q|A'|$. Since at each recursive step, the size of the graph with which the algorithm calls itself decreases, eventually either A' becomes empty or $B \setminus N_G(A \setminus A')$ becomes empty. Hence, the algorithm outputs no. Now we need to show that if there exist sets (A^*, B^*) such that there is a q -expansion of A^* into B^* , then at each recursive call, we have that $A^* \subseteq A$ and $B^* \subseteq B$. At the start of the algorithm, $A^* \subseteq A$ and $B^* \subseteq B$. Since $N(B^*) \subseteq A^*$ and for all $u \in A^* d_G(u) \geq q$, we have that $A^* \cup B^* \subseteq V(G')$. If $N(B') \subseteq A'$, then the algorithm of Lemma 6 when run on G', q will output (A^*, B^*) . Note that $B^* \subseteq B'$. At the recursive step, $A^* \subseteq A'$ and since $B^* \cap N_G(A \setminus A') = \emptyset$, we have that $B^* \subseteq B' \setminus N_G(A \setminus A')$. Hence, $G[A^* \cup B^*]$ is a subgraph of $G[A' \cup (B \setminus N_G(A \setminus A'))]$ which concludes the correctness of the algorithm. Since at each recursive call the size of the graph decreases by at least 1, the total time taken by the above algorithm is polynomial in n . ◀

One may think of a q -expansion in a bipartite graph with bipartition (A, B) as an allocation of the items in B to each customer in A such that every customer gets at least q items. For our kernel we will need a generalization of q -expansions to the setting where the items in B have different values, and every customer gets items of total value at least q .

► **Definition 8** (Weighted q -expansion). Let $G := ((A, B), E)$ be a bipartite graph with capacity function $w_b : B \rightarrow \mathbb{N}$. Then, a weighted q -expansion in G is an edge weight function $f : E(G) \rightarrow \mathbb{N}$ that satisfies the capacity constraints w_b and also satisfies the demand constraints $w_a = q$. For an integer $W \in \mathbb{N}$, the q -expansion f is called a W -strict q -expansion if f allocates at least $q + W - 1$ value to at least one vertex r in A , and in this case we say that f is W -strict at r . Further, a q -expansion f is *strict* (at r) if it is 1-strict (at r). If f is unsplitting we call f an *unsplitting* q -expansion.

► **Lemma 9.** *There exists a polynomial-time algorithm that given a bipartite graph $G := ((A, B), E)$, an integer q and a capacity function $w_b : B \rightarrow \mathbb{N}$ outputs (if it exist) two sets $X \subseteq A$ and $Y \subseteq B$ along with a weighted q -expansion in $G[X \cup Y]$ such that $N(Y) \subseteq X$.*

Proof. Construct the twin graph $T_{AB} := (A, B')$ of G . Run the algorithm of Lemma 7 with input T_{AB}, q that outputs sets $X \subseteq A$ and $Y' \subseteq B'$ such that X has q -expansion into Y' and $N(Y') \subseteq X$. Consider the set $Y := \{v \in B \mid v_i \in Y'\}$. Define a weight function $f : E(G[X \cup Y]) \rightarrow \mathbb{N}$ as follows: for all $uv \in E(G[X \cup Y])$ $f(uv) = |\{v_i \in Y' \mid v_i \text{ matched to } u\}|$.

Clearly, $N(Y) \subseteq X$. Now we claim that f is a weighted q -expansion in $G[X \cup Y]$ with capacity function w_b and demand function $w_a = q$. For any vertex $u \in A$, there are at least q vertices in Y' are matched to u . Hence for all $u \in A$, we have that $\sum_{v \in N(u)} f(uv) \geq q = w_a$. At the same time, for any vertex $v \in B$, there are at most $w_b(v)$ copies of v in Y' . Therefore, for all $v \in Y$ we have $\sum_{u \in N(v)} f(uv) \leq w_b(v)$. ◀

► **Lemma 10.** *There exists a polynomial-time algorithm that given a weighted q -expansion $f : E(G) \rightarrow \mathbb{N}$ in $G := ((A, B), E)$, a capacity function $w_b : B \rightarrow \mathbb{N}$ and an integer W such that $W = \max_{e \in E(G)} f(e)$ outputs an unsplitting W -strict weighted $(q - W + 1)$ -expansion in G .*

Proof. Run the algorithm of Lemma 2 with inputs $G, f, w_a = q, w_b, W$ and a vertex $u \in A$. In case f is strict, u is the vertex r that makes f strict. Let the function $h : E(G) \rightarrow \mathbb{N}$ be the output of Lemma 2. Now h is an unsplitting edge weight function that satisfies the capacity constraints, satisfies the demands $q - W + 1$, and additionally satisfies the demand q of u . Hence, h is the required unsplitting weighted W -strict $(q - W + 1)$ -expansion in G . ◀

► **Lemma 11** (Weighted Expansion Lemma). *Let $q, W \geq 1$ be positive integers and G be a bipartite graph with vertex bipartition (A, B) and $w_b : B \rightarrow \{1, \dots, W\}$ be a capacity function such that $\sum_{v \in B} w_b(v) \geq (q + W - 1) \cdot |A|$, and there are no isolated vertices in B . Then there exist nonempty vertex sets $X \subseteq A$ and $Y \subseteq B$ such that $N(Y) \subseteq X$ and there is an unsplitting weighted W -strict q -expansion of X into Y . Furthermore, the sets X and Y can be found in time polynomial in the size of G .*

Proof. Construct the twin graph T_{AB} from G and w_b , the bipartition of T_{AB} is (A, B') . Now, obtain using the Expansion Lemma 6 with $q' = q + W - 1$ on T_{AB} sets $X \subseteq A$ and $Y' \subseteq B'$, such that $N(Y') \subseteq X$ and there is a $(q + W - 1)$ -expansion from X to Y' in T_{AB} .

Let $Y := \{v \in B \mid v_i \in Y'\}$ (here the $v_i \in Y'$ are as in Definition 4). Then $N(Y) \subseteq X$ and the $(q + W - 1)$ -expansion from X to Y' in T_{AB} immediately yields a weighted $(q + W - 1)$ -expansion f from X to Y in G . Applying Lemma 10 on $G[X \cup Y]$ using the weighted $(q + W - 1)$ -expansion f proves the statement of the lemma. ◀

5

 Obtaining the Linear Kernel

► **Definition 12.** For a graph G and a pair of vertex-disjoint sets $X, Y \subseteq V(G)$, we define the weighted graph \tilde{G}_{XY} as follows: $V(\tilde{G}_{XY}) := X \cup \tilde{Y}$ such that there is a bijection $h : cc(G[Y]) \rightarrow \tilde{Y}$ where $cc(G[Y])$ is the set of connected components of $G[Y]$. $E(\tilde{G}_{XY}) := \{xc \mid x \in X, c \in \tilde{Y}, c = h(C) \text{ and } x \in N_G(C)\}$. We also define a weight function $w : \tilde{Y} \rightarrow \mathbb{N}$ such that for all $c \in \tilde{Y}$, $w(c) = |h^{-1}(c)|$.

► **Definition 13 (Reducible Pair).** For a graph G , a pair of vertex-disjoint sets (X, Y) where $X, Y \subseteq V(G)$ is called a (strict) reducible pair if $N(Y) \subseteq X$, the size of every component in $G[Y]$ is at most ℓ , and there exists a (strict) weighted $(2\ell - 1)$ -expansion in \tilde{G}_{XY} .

► **Definition 14.** A reducible pair (X, Y) is called minimal if there is no reducible pair (X', Y') such that $X' \subset X$ and $Y' \subseteq Y$.

► **Lemma 15.** *There exists a polynomial-time algorithm that given an ℓ -COC instance (G, k) together with a vertex-disjoint set pair $A, B \subseteq V(G)$ outputs (if it exists) a reducible pair (X, Y) where $X \subseteq A$ and $Y \subseteq B$.*

Proof. Construct $\tilde{G}_{AB} := (A, \tilde{B})$ and run the algorithm of Lemma 9 with input $\tilde{G}_{AB}, w, q = 2\ell - 1$ which outputs sets $X \subseteq A$ and $Y' \subseteq \tilde{B}$ (if it exists) along with a weighted $(2\ell - 1)$ -expansion of X into Y' such that $N(Y') \subseteq X$. Now from Y' we obtain the set $Y := \bigcup_{y \in Y'} h^{-1}(y)$. Clearly, $N(Y) \subseteq X$ and hence, (X, Y) is the desired reducible pair. ◀

► **Lemma 16.** *Given an ℓ -COC instance (G, k) , if $|V(G)| \geq 2lk$ and (G, k) is a YES-instance, then there exists a reducible pair (X, Y) .*

Proof. Without loss of generality, we can assume that G is a connected graph. Let S be an ℓ -COC solution of size at most k . Clearly, $|V \setminus S| \geq (2\ell - 1)k$. We define $A := S$ and $B := V \setminus S$ and construct $\tilde{G}_{AB} = (A, \tilde{B})$. We have the weight function $w_b : \tilde{B} \rightarrow \mathbb{N}$ such that for all $v \in \tilde{B}$, $w_b(v) = |h^{-1}(v)| \leq \ell$, as the size of components in $G[V \setminus S]$ is at most ℓ . We have that $\sum_{v \in \tilde{B}} w_b(v) \geq (2\ell - 1)|A|$ and there are no isolated vertices in \tilde{B} . Hence, (A, B) is the desired reducible pair. ◀

► **Lemma 17.** *Let (X, Y) be a reducible pair. Then, there exists a partition of $X \cup Y$ into $C_1, \dots, C_{|X|}$ such that (i) for all $u_i \in X$, we have $u_i \in C_j$ if and only if $i = j$, (ii) for all $i \in [|X|]$, $|C_i| \geq \ell + 1$, (iii) for every component C in $G[Y]$, there exists a unique C_i such that $V(C) \subseteq C_i$ and $u_i \in N(C)$ and (iv) if (X, Y) is a strict reducible pair, then there exists C_j such that $|C_j| \geq 2\ell + 1$.*

Proof. Construct $\tilde{G}_{XY} := (X, \tilde{Y})$. Run the algorithm of Lemma 10 with input $\tilde{G}_{XY}, q = 2\ell - 1$, and $W = \ell$ (as the capacity of any vertex in \tilde{Y} is at most ℓ) which outputs an unsplitting weighted ℓ -expansion f' in \tilde{G}_{XY} . In polynomial time, we modify f' such that if there is a vertex $v \in \tilde{Y}$ such that $\forall u \in N(v), f'(uv) = 0$, we choose a vertex $u \in N(v)$ and set $f'(uv) = w_b(v)$. For each $u_i \in X$ define $C_i := u_i \bigcup_{f'(u_i v) \neq 0} h^{-1}(v)$. Since f' is unsplitting, the collection $C_1, \dots, C_{|X|}$ forms a partition of $X \cup Y$. By the definition of C_i , we have that for any $u_i \in X$, $u_i \in C_j$ if and only if $i = j$. For any component C in $G[Y]$, $h(C)$ is matched to a unique vertex $u_i \in X$ by f' , we have that $V(C) \subseteq C_i$. As f' is a weighted ℓ -expansion, $|C_i| = 1 + \sum_{f'(u_i v) \neq 0} |h^{-1}(v)| = 1 + \sum_{f'(u_i v) \neq 0} f'(u_i v) \geq 1 + \ell$. Let (X, Y) be strict at $u_j \in X$. Then, we can use Lemma 10 to obtain the expansion f' such that it is strict at u_j . Hence, $|C_j| = 1 + \sum_{f'(u_j v) \neq 0} |h^{-1}(v)| = 1 + \sum_{f'(u_j v) \neq 0} f'(u_j v) > 1 + \ell + (\ell - 1)$ which implies $|C_j| \geq 2\ell + 1$. This concludes the proof of the lemma. ◀

► **Lemma 18.** *Let (X, Y) be a reducible pair. If (G, k) is a YES-instance for ℓ -COC, then there exists an ℓ -COC solution S of size at most k such that $X \subseteq S$ and $S \cap Y = \emptyset$.*

Proof. By Lemma 17 we have that there are $C_1, \dots, C_{|X|} \subseteq X \cup Y$ vertex disjoint sets of size at least $\ell + 1$ such that for all $i \in [|X|]$, $G[C_i]$ is a connected set. Let S' be an arbitrary solution. Then, S' must contain at least one vertex from each C_i . Let $S := S' \setminus (X \cup Y) \cup X$. We have that $|S| \leq |S'| - |X| + |X| = |S'|$. As any connected set of size $\ell + 1$ that contains a vertex in Y also contains a vertex in X and $X \subseteq S$, S is also an ℓ -COC solution. ◀

Now we encode an ℓ -COC instance (G, k) as an INTEGER LINEAR PROGRAMMING instance. We introduce $n = |V(G)|$ variables, one variable x_v for each vertex $v \in V(G)$. Setting the variable x_v to 1 means that v is in S , while setting $x_v = 0$ means that v is not in S . To ensure that S contains a vertex from every connected set of size $\ell + 1$, we can introduce constraints $\sum_{v \in C} x_v \geq 1$ where C is a connected set of size $\ell + 1$. The size of S is given by $\sum_{v \in V(G)} x_v$. This gives us the following ILP formulation:

$$\begin{aligned} & \text{minimize} && \sum_{v \in V(G)} x_v, \\ & \text{subject to} && \sum_{v \in C} x_v \geq 1 \quad \text{for every connected set } C \text{ of size } \ell + 1 \\ & && 0 \leq x_v \leq 1 \quad \text{for every } v \in V(G) \\ & && x_v \in \mathbb{Z} \quad \text{for every } v \in V(G). \end{aligned}$$

Note that there are $n^{\mathcal{O}(\ell)}$ connected sets of size at most ℓ in a graph on n vertices. Hence, providing an explicit ILP requires $n^{\mathcal{O}(\ell)}$ time which forms the bottleneck for the runtime of the kernelization algorithm that follows. We consider the Linear Programming relaxation of above ILP obtained by dropping the constraint that $x \in \mathbb{Z}$. By an optimal LP solution S_L with weight L we mean the set of values assigned to each variable, and optimal value is L . For a set of vertices $X \subseteq V(G)$, $X = 1$ ($X = 0$) denotes that every variable corresponding to vertices in X is set to 1 (0).

► **Lemma 19.** *Let S_L be an optimal LP solution for G such that $x_v = 1$ for some $v \subseteq V(G)$. Then, $S_L - x_v$ is an optimal LP solution for $G - v$ of value $L - 1$.*

Proof. Clearly, $S_L - x_v$ is feasible solution for $G - v$ of value $L - 1$. Suppose it is not optimal. Let $S_{L'}$ be an optimal LP solution for $G - v$ such that $L' < L - 1$. Then, $S_{L'} \cup x_v$ with $x_v = 1$ is an optimal LP solution for G with value $< L - 1 + 1 = L$ contradicting that the optimal solution value of LP for G is L . ◀

From now on by running LP after setting $x_v = 1$ for some vertex v , we mean running the LP algorithm for $G - v$ and including $x_v = 1$ in the obtained solution to get a solution for G .

► **Lemma 20.** *Let (X, Y) be a strict reducible pair. Then every optimal LP solution sets at least one variable corresponding to a vertex in X to 1.*

Proof. By Lemma 18, we have that every connected set of size $\ell + 1$ in $G[X \cup Y]$ contains a vertex in X . Hence, from any LP solution S_L , a feasible LP solution can be obtained by setting $X = 1$ and $Y = 0$. Since, we have at least $|X|$ many vertex disjoint LP constraints, for each $v_i \in X$, we have $\sum_{u \in C_i} x_u = 1$. By Lemma 17, there is a set $C_j \subseteq X \cup Y$ such that $|C_j| \geq 2\ell + 1$. If $x_{v_j} \neq 1$, then there is a vertex $w \in C_j$ such that $x_w > 0$. Let $w \in C \subseteq C_j$ where $G[C]$ is a connected component in $G[Y]$. Since $|C| \leq \ell$, there is a connected set C' of size at least $\ell + 1$ in $G[C_j] - C$. But now $\sum_{u \in C'} x_u < 1$ contradicting that S_L is feasible. ◀

► **Lemma 21.** *Let (X, Y) be a minimal reducible pair. If for any vertex $v \in X$, an optimal LP solution sets $x_v = 1$, then it also sets $X = 1$ and $Y = 0$.*

Proof. We prove the lemma by contradiction. Let $X' \subset X$ be the largest subset of X such that $X' = 1$. Consider \tilde{G}_{XY} . Let $Y' \subseteq \tilde{Y}$ be the set of vertices such that $N(Y') \subseteq X'$. Let $Z := \bigcup_{v \in Y'} h^{-1}(v)$. By the minimality of (X, Y) , we have that $\sum_{v \in Y'} w(v) < (2\ell - 1)|X'|$. Hence, $\sum_{v \in \tilde{Y} \setminus Y'} w(v) > (2\ell - 1)|X \setminus X'|$. Clearly, the weighted $(2\ell - 1)$ -expansion in the reducible pair (X, Y) when restricted to $(X \setminus X', Y \setminus Z)$ provides a weighted $(2\ell - 1)$ -expansion of $X \setminus X'$ into $Y \setminus Z$. This implies that $(X \setminus X', Y \setminus Z)$ is a strict reducible pair in $G - (X' \cup Z)$. By Lemma 19, we have that the LP solution restricted to $G - (X' \cup Z)$ is optimal. Since $(X \setminus X', Y \setminus Z)$ is a strict reducible pair, by Lemma 20, there is a vertex $u \in X \setminus X'$ such that $x_u = 1$, but this contradicts the maximality of X' . Therefore, if for any vertex $v \in X$, an LP solution sets $x_v = 1$, then it sets $X = 1$ and $Y = 0$. ◀

► **Lemma 22.** *There exists a polynomial time algorithm that given an integer ℓ and ℓ -COC instance (G, k) on at least $2\ell k$ vertices either finds a reducible pair (X, Y) or concludes that (G, k) is a NO-instance.*

Proof. If (G, k) is a YES-instance of ℓ -COC, then by Lemma 16, there exists a reducible pair (X, Y) . We use the following algorithm to find one:

Step 1. Run the LP algorithm. Let $A = 1$ and $B = 0$ in the LP solution.

Step 2. If both A and B are non-empty, then run the algorithm of Lemma 15 with input $(G, k), A, B$. If it outputs a reducible pair (X, Y) , then return (X, Y) and terminate. Otherwise, go to step 3.

Step 3. Now we do a linear search for a vertex in X . For each vertex $v \in V(G)$, do the following: in the original LP introduce an additional constraint that sets the value of the variable x_v to 1 i.e. $x_v = 1$ and run the LP algorithm. If the optimal value of the new LP is the same as the optimal value of the original LP, then let $A = 1$ and $B = 0$ be the sets of variables set to 1 and 0 respectively in the optimal solution of the new LP and go to step 2.

Step 4. Output a trivial NO-instance.

Step 1 identifies the set of variables set to 1 and 0 by the LP algorithm. By Lemma 21, we have that if there is a minimal reducible pair (X, Y) in G , then $X \subseteq A$ and $Y \subseteq B$. So, in Step 2 if the algorithm succeeds in finding one, we return the reducible pair and terminate otherwise we look for a potential vertex in X and set it to 1. If (X, Y) exists, then for at least one vertex, setting $x_v = 1$ would set $X = 1$ and $Y = 0$ (by Lemma 21) without changing the LP value and we go to Step 2 to find it. If for each choice of $v \in V(G)$, the LP value changes when x_v is set to 1, we can conclude that there is no reducible pair and output a trivial NO instance. Since, we need to do this search at most n times and each step takes only polynomial time, the total time taken by the algorithm is polynomial in the input size. ◀

► **Theorem 23.** *For every constant $\ell \in \mathbb{N}$, ℓ -COMPONENT ORDER CONNECTIVITY admits a kernel with at most $2\ell k$ vertices that takes $n^{\mathcal{O}(\ell)}$ time.*

References

- 1 Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *SIGecom Exchanges*, 5(3):11–18, 2005.
- 2 Maw-Shang Chang, Li-Hsuan Chen, Ling-Ju Hung, Peter Rossmanith, and Ping-Chen Su. Fixed-parameter algorithms for vertex cover p3. *Discrete Optimization*, 19:12–22, 2016. doi:10.1016/j.disopt.2015.11.003.
- 3 Jianer Chen, Iyad A. Kanj, and Weijia Jia. Vertex cover: Further observations and further improvements. *J. Algorithms*, 41(2):280–301, 2001. doi:10.1006/jagm.2001.1186.

- 4 Jianer Chen, Iyad A. Kanj, and Ge Xia. Improved upper bounds for vertex cover. *Theor. Comput. Sci.*, 411(40-42):3736–3756, 2010.
- 5 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 6 Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of mathematics*, pages 439–485, 2005.
- 7 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
- 8 Pål Grønås Drange, Markus Sortland Dregi, and Pim van 't Hof. On the computational complexity of vertex integrity and component order connectivity. In *Algorithms and Computation – 25th International Symposium, ISAAC 2014, Jeonju, Korea, December 15-17, 2014, Proceedings*, pages 285–297, 2014. doi:10.1007/978-3-319-13075-0_23.
- 9 Fedor V. Fomin, Serge Gaspers, Dieter Kratsch, Mathieu Liedloff, and Saket Saurabh. Iterative compression and exact algorithms. *Theor. Comput. Sci.*, 411(7-9):1045–1053, 2010.
- 10 Fedor V. Fomin, Fabrizio Grandoni, and Dieter Kratsch. A measure & conquer approach for the analysis of exact algorithms. *J. ACM*, 56(5), 2009.
- 11 Fedor V. Fomin and Dieter Kratsch. *Exact Exponential Algorithms*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2010.
- 12 D. Gross, M. Heinig, L. Iswara, W. Kazmierczak, K. Luttrell, J. T. Saccoman, and C. Suffel. A survey of component order connectivity models of graph theoretic networks. *WSEAS Transactions on Mathematics*, 12:895–910, 2013.
- 13 Frantisek Kardos, Ján Katrenic, and Ingo Schiermeyer. On computing the minimum 3-path vertex cover and dissociation number of graphs. *Theor. Comput. Sci.*, 412(50):7009–7017, 2011. doi:10.1016/j.tcs.2011.09.009.
- 14 Stefan Kratsch. Recent developments in kernelization: A survey. *Bulletin of the EATCS*, 113, 2014.
- 15 Jan Karel Lenstra, David B. Shmoys, and Éva Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Math. Program.*, 46:259–271, 1990.
- 16 John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is np-complete. *J. Comput. Syst. Sci.*, 20(2):219–230, 1980. doi:10.1016/0022-0000(80)90060-4.
- 17 Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Kernelization–preprocessing with a guarantee. In *The Multivariate Algorithmic Revolution and Beyond*, pages 129–161. Springer, 2012.
- 18 George L. Nemhauser and Leslie E. Trotter Jr. Properties of vertex packing and independence system polyhedra. *Math. Program.*, 6(1):48–61, 1974. doi:10.1007/BF01580222.
- 19 J. M. Robson. Algorithms for maximum independent sets. *J. Algorithms*, 7(3):425–440, 1986.
- 20 Jianhua Tu. A fixed-parameter algorithm for the vertex cover p_3 problem. *Inf. Process. Lett.*, 115(2):96–99, 2015.
- 21 Jianhua Tu and Wenli Zhou. A factor 2 approximation algorithm for the vertex cover p_3 problem. *Inf. Process. Lett.*, 111(14):683–686, July 2011. doi:10.1016/j.ipl.2011.04.009.
- 22 David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.
- 23 Mingyu Xiao and Shaowei Kou. Faster computation of the maximum dissociation set and minimum 3-path vertex cover in graphs. In *Frontiers in Algorithmics – 9th International Workshop, FAW 2015, Guilin, China, July 3-5, 2015, Proceedings*, pages 282–293, 2015. doi:10.1007/978-3-319-19647-3_26.

20:14 **A $2\ell k$ Kernel for ℓ -Component Order Connectivity**

- 24 Mingyu Xiao and Hiroshi Nagamochi. Exact algorithms for maximum independent set. In *Algorithms and Computation – 24th International Symposium, ISAAC 2013, Hong Kong, China, December 16-18, 2013, Proceedings*, pages 328–338, 2013.
- 25 Mihalis Yannakakis. Node-deletion problems on bipartite graphs. *SIAM J. Comput.*, 10(2):310–327, 1981. doi:10.1137/0210022.