Fixed parameter complexity of distance constrained labeling and uniform channel assignment problems *

Jiří Fiala   Tomáš Gavenčiak   Dušan Knop
Martin Koutecký   Jan Kratochvíl[†]

Department of Applied Mathematics, Charles University
Malostranské nám. 25, Prague, Czech Republic
{fiala,gavento,knop,koutecky,honza}@kam.mff.cuni.cz

Abstract

We study computational complexity of the class of distance-constrained graph labeling problems from the fixed parameter tractability point of view. The parameters studied are neighborhood diversity and clique width.

We rephrase the distance constrained graph labeling problem as a specific uniform variant of the CHANNEL ASSIGNMENT problem and show that this problem is fixed parameter tractable when parameterized by the neighborhood diversity together with the largest weight. Consequently, every $L(p_1, p_2, \ldots, p_k)$-LABELING problem is FPT when parameterized by the neighborhood diversity, the maximum p_i and k.

Our results yield also FPT algorithms for all $L(p_1, p_2, \ldots, p_k)$-LABELING problems when parameterized by the size of a minimum vertex cover, answering an open question of Fiala et al.: *Parameterized complexity of coloring problems: Treewidth versus vertex cover.* The same consequence applies on CHANNEL ASSIGNMENT when the maximum weight is additionally included among the parameters.

*Paper supported by project Kontakt LH12095 and by GAUK project 1784214. Extended abstract of this paper was accepted to the conference COCOON 2016

[†]Second, third and fourth authors are supported by the project SVV–2016–260332. First, third and fifth authors are supported by project CE-ITI P202/12/G061 of GA ČR.

Finally, we show that the uniform variant of the CHANNEL ASSIGNMENT problem becomes NP-complete when generalized to graphs of bounded clique width.

1 Introduction

The frequency assignment problem in wireless networks yields an abundance of various mathematical models and related problems. We study a group of such discrete optimization problems in terms of parameterized computational complexity, which is one of the central paradigms of contemporary theoretical computer science. We study parameterization of the problems by *clique width* and particularly by *neighborhood diversity* (nd), a graph parameter lying between clique width and the size of a minimum vertex cover.

All these problems are NP-hard even for constant clique width, including the uniform variant, as we show in this paper. On the other hand, we prove that they are in FPT with respect to nd. Such fixed parameter tractability has so far only been known only for the special case of $L(p, 1)$ labeling when parameterized by vertex cover [7].

1.1 Distance constrained labelings

Given a k-tuple of positive integers $p_1, \dots, p_k$, called *distance constraints*, an $L(p_1, \dots, p_k)$-labeling of a graph is an assignment l of integer labels to the vertices of the graph satisfying the following condition: Whenever vertices u and v are at distance i, the assigned labels differ by at least p_i. Formally, $\mathrm{dist}(u, v) = i \implies |l(u) - l(v)| \geq p_i$ for all $u, v : \mathrm{dist}(u, v) \leq k$. Often only non-increasing sequences of distance constraints are considered.

Any $L(1)$-labeling is a graph coloring and vice-versa. Analogously, any coloring of the k-th distance power of a graph is an $L(1, \dots, 1)$-labeling. The concept of $L(2, 1)$-labeling is attributed to Roberts by Griggs and Yeh [13]. It is not difficult to show that whenever l is an optimal $L(p_1, \dots, p_k)$-labeling within a range $[0, \lambda]$, then the so called *span* λ is a linear combination of $p_1, \dots, p_k$ [13, 16]. In particular, a graph G allows an $L(p_1, \dots, p_k)$-labeling of span λ if and only if it has an $L(cp_1, \dots, cp_k)$-labeling of span $c\lambda$ for any positive integer c.

For computational complexity purposes, we define the following class of decision problems:

Problem 1. $L(p_1, \dots, p_k)$-LABELING:

Parameters:	Positive integers $p_1, \dots, p_k$.
Input:	Graph G, positive integer λ.
Query:	Is there an $L(p_1, \dots, p_k)$ labeling of G using labels from the interval $[0, \lambda]$?

The $L(2,1)$-LABELING problem was shown to be NP-complete by Griggs and Yeh [13] by a reduction from HAMILTONIAN CYCLE (with $\lambda = |V_G|$). Fiala, Kratochvíl and Kloks [8] showed that $L(2,1)$-LABELING remains NP-complete also for all fixed $\lambda \geq 4$, while for $\lambda \leq 3$ it is solvable in linear time.

Despite a conjecture that $L(2,1)$-LABELING remains NP-complete on trees [13], Chang and Kuo [2] showed a dynamic programming algorithm for this problem, as well as for all $L(p_1, p_2)$-labelings where p_2 divides p_1. All the remaining cases of the $L(p_1, p_2)$-LABELING problem on trees have been shown to be NP-complete by Fiala, Golovach and Kratochvíl [6]. The same authors showed that $L(2,1)$-LABELING is already NP-complete on series-parallel graphs [5], which have of tree width at most 2. Note that these results imply NP-hardness of $L(3,2)$-LABELING on graphs of clique width at most 3 and of $L(2,1)$-LABELING for clique width at most 6 [3].

On the other hand, when λ is fixed, then the existence of an $L(p_1, \dots, p_k)$-labeling of G can be expressed in MSO_1, hence it allows a linear time algorithm on any graph of bounded clique width [15].

Fiala et al. [7] showed that the problem of $L(p,1)$-LABELING is FPT when parameterized by p together with the size of the vertex cover. They also ask for the complexity characterization of the related CHANNEL ASSIGNMENT problem. We extend their work to the broader class of graphs and, consequently, in our Theorem 9 we provide a solution for their open problem.

1.2 Channel assignment

Channel assignment is a concept closely related to distance constrained graph labeling. Here, every edge has a prescribed weight $w(e)$ and it is required that the labels of adjacent vertices differ at least by the weight of the corresponding edge. The associated decision problem is defined as follows:

Problem 2. CHANNEL ASSIGNMENT:

Input:	Graph G, a positive integer λ, edge weights $w : E_G \to \mathbb{N}$.
Query:	Is there a labeling l of the vertices of G by integers from $[0, \lambda]$ such that $|l(u) - l(v)| \geq w(u, v)$ for all $(u, v) \in E_G$?

The maximal edge weight is an obvious necessary lower bound for the span of any labeling. Observe that for any bipartite graph, in particular also for all trees, it is also an upper bound — a labeling that assigns 0 to one class of the bipartition and $w_{\max} = \max\{w(e), e \in E_G\}$ to the other class satisfies all edge constraints. McDiarmid and Reed [19] showed that it is NP-complete to decide whether a graph of tree width 3 allows a channel assignment of given span λ. This NP-hardness hence applies on graphs of clique width at most 12 [3]. It is worth noting that for graphs of tree width 2, i.e. for subgraphs of series-parallel graphs, the complexity characterization of CHANNEL ASSIGNMENT is still open. Only a few partial results are known [20], among others that CHANNEL ASSIGNMENT is polynomially solvable on graphs of bounded tree width if the span λ is bounded by a constant.

Any instance G, λ of the $L(p_1, \ldots, p_k)$-LABELING problem can straightforwardly be reduced to an instance G^k, λ, w of the CHANNEL ASSIGNMENT problem. Here, G^k arises from G by connecting all pairs of vertices that are in G at distance at most k, and for the edges of G^k we let $w(u, v) = p_i$ whenever $\text{dist}_G(u, v) = i$.

The resulting instances of CHANNEL ASSIGNMENT have by the construction some special properties. We explore and generalize these to obtain a uniform variant of the CHANNEL ASSIGNMENT problem.

1.3 Neighborhood diversity

Lampis significantly reduced (from the tower function to double exponential) the hidden constants of the generic polynomial algorithms for MSO_2 model checking on graphs with bounded vertex cover [17]. To extend this approach to a broader class of graphs he introduced a new graph parameter called the neighborhood diversity of a graph as follows:

Definition 3 (Neighborhood diversity). *A partition $V_1, \ldots, V_d$ is called a* neighborhood diversity decomposition *if it satisfies*

- *each V_i induces either an empty subgraph or a complete subgraph of G, and*

- *for each distinct V_i and V_j there are either no edges between V_i and V_j, or every vertex of V_i is adjacent to all vertices of V_j.*

We write $u \sim v$ to indicate that u and v belong to the same class of the decomposition.

The neighborhood diversity *of a graph G, denoted by* $\mathrm{nd}(G)$*, is the minimum τ such that G has a neighborhood diversity decomposition with τ classes.*

Observe that for the optimal neighborhood diversity decomposition it holds that $u \sim u'$ is equivalent with $N(u) \setminus v = N(v) \setminus u$. Therefore, the optimal neighborhood diversity decomposition can be computed in $O(n^3)$ time [17].

Classes of graphs of bounded neighborhood diversity reside between classes of bounded vertex cover and graphs of bounded clique width. Several non-MSO_1 problems, e.g. HAMILTONIAN CYCLE can be solved in polynomial time on graphs of bounded clique width [21]. On the other hand, Fomin et al. stated more precisely that the HAMILTONIAN CYCLE problem is $W[1]$-hard, when parameterized by clique width [9]. In sequel, Lampis showed that some of these problems, including HAMILTONIAN CYCLE, are indeed fixed parameter tractable on graphs of bounded neighborhood diversity [17].

Ganian and Obdržálek [12] further deepened Lampis' results and showed that also problems expressible in MSO_1 with cardinality constraints ($\mathrm{cardMSO}_1$) are fixed parameter tractable when parameterized by $\mathrm{vc}(G)$ or by $\mathrm{nd}(G)$.

Any vertex cover U can be converted into a neighborhood diversity decomposition with at most $2^{|U|} + |U|$ classes. Hence any graph G satisfies $\mathrm{nd}(G) \leq 2^{\mathrm{vc}(G)} + \mathrm{vc}(G)$ where $\mathrm{vc}(G)$ is the size of minimal vertex cover of G. This construction is used in the proof of Theorem 9 (see section 4 for more details).

Observe that a sufficiently large n-vertex graph of bounded neighborhood diversity can be described in significantly more effective way, namely by using only $O(\log n \, \mathrm{nd}(G)^2)$ space:

Definition 4 (Type graph). *The* type graph $T(G)$ *for a neighborhood diversity decomposition $V_1, \ldots, V_d$ of a graph G is a vertex weighted graph on vertices $\{t_1, \ldots, t_d\}$, where each t_i is assigned weight $s(t_i) = |V_i|$, i.e. the size of the corresponding class of the decomposition. Distinct vertices t_i and t_j are adjacent in $T(G)$ if and only if the edges between the two corresponding classes V_i and V_j form a complete bipartite graph. Moreover, $T(G)$ contains* *a loop incident with vertex t_i if and only if the corresponding class V_i induces a clique.*

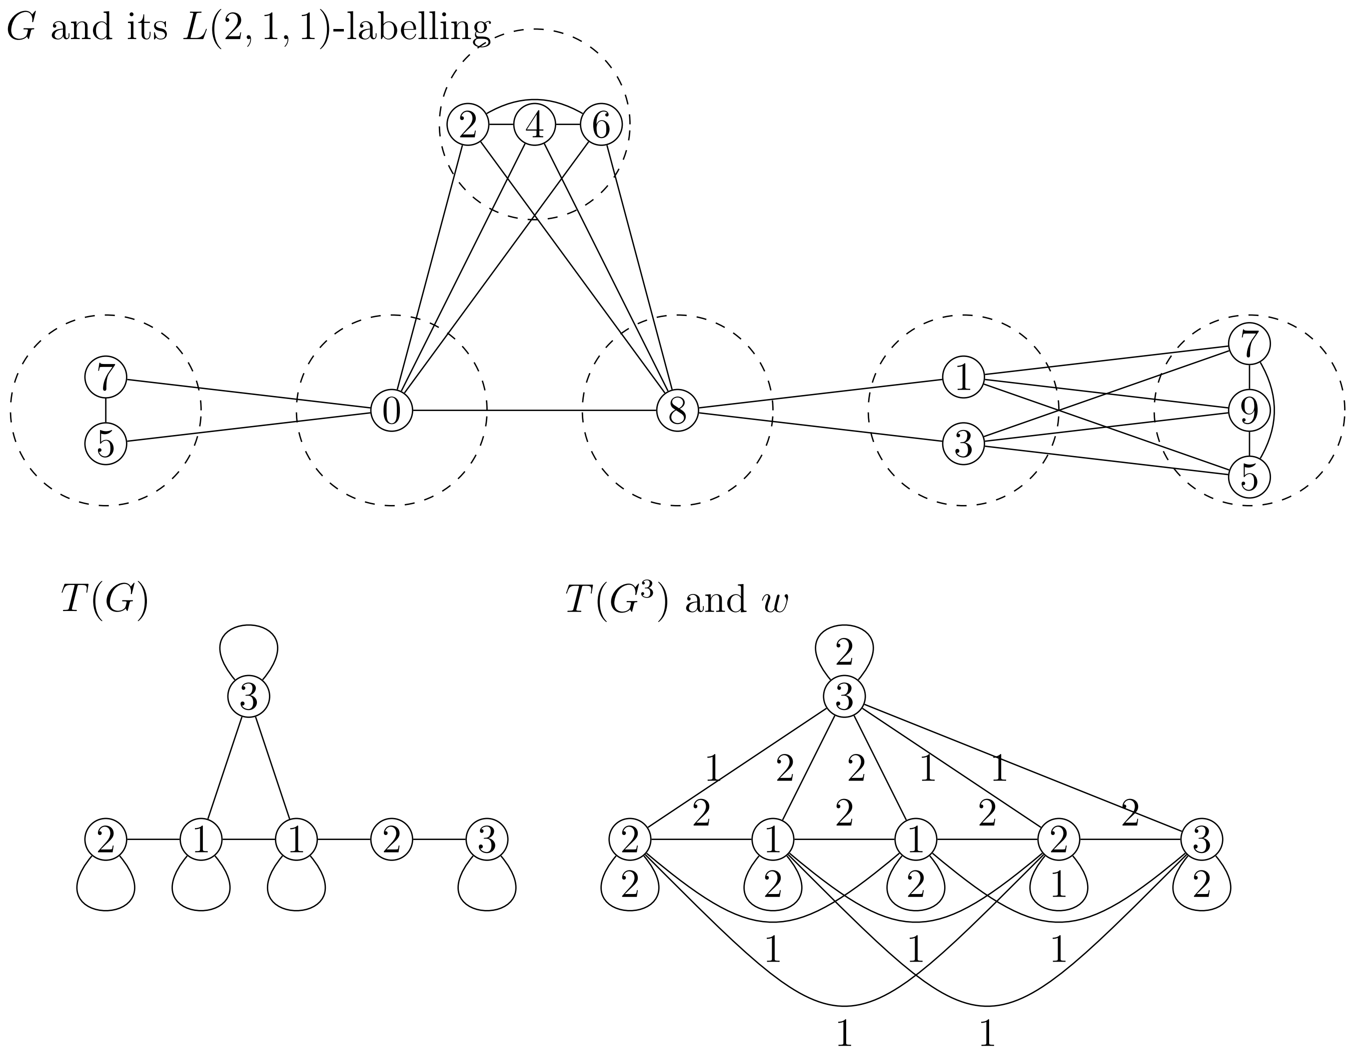

Figure 1: An example of a graph with its neighborhood diversity decomposition. Vertex labels indicate one of its optimal $L(2, 1, 1)$-labelings. The corresponding type graph. The weighted type graph corresponding to the resulting instance of the Channel Assignment problem.

For our purposes, i.e. to decide existence of a suitable labeling of a graph G, it suffices to consider only its type graph, as G can be uniquely reconstructed from $T(G)$ (upto an isomorphism) and vice-versa.

Moreover, the reduction of $L(p_1, \ldots, p_k)$-labeling to Channel Assignment preserves the property of bounded neighborhood diversity:

Observation 5. *For any graph G and any positive integer k it holds that* $\mathrm{nd}(G) \geq \mathrm{nd}(G^k)$.

Proof. The optimal neighborhood diversity decomposition of G is a neighborhood diversity decomposition of G^k. □

1.4 Our contribution

Our goal is an extension of the FPT algorithm for $L(2,1)$-LABELING on graphs of bounded vertex cover to broader graph classes and for rich collections of distance constraints. In particular, we aim at $L(p_1, \ldots, p_k)$-LABELING on graphs of bounded neighborhood diversity.

For this purpose we utilize the aforementioned reduction to CHANNEL ASSIGNMENT, taking into account that the neighborhood diversity remains bounded, even though the underlying graph changes.

It is worth to note that we must adopt additional assumptions for the CHANNEL ASSIGNMENT since otherwise it is NP-complete already on complete graphs, i.e. on graphs with $\mathrm{nd}(G) = 1$. To see this, we recall the construction of Griggs and Yeh [13]. They show that a graph H on n vertices has a Hamiltonian path if and only if the complement of H extended by a single universal vertex allows an $L(2,1)$-labeling of span $n+1$. As the existence of a universal vertex yields diameter at most two, the underlying graph for the resulting instance of CHANNEL ASSIGNMENT is K_{n+1}.

On the other hand, the additional assumptions on the instances of CHANNEL ASSIGNMENT will still allow us to reduce any instance of the $L(p_1, \ldots, p_k)$-LABELING problem. By the reduction, all edges between classes of the neighborhood diversity decomposition are assigned the same weight. We formally adopt this as our additional constraint as follows:

Definition 6. *The edge weights w on a graph G are* nd-uniform *if $w(u,v) = w(u',v')$ whenever $u \sim u'$ and $v \sim v'$ with respect to the optimal neighborhood diversity decomposition. In a similar way we define uniform weights with respect to a particular decomposition.*

Our main contribution is an algorithm for the following scenario:

Theorem 7. *The* CHANNEL ASSIGNMENT *problem on* nd*-uniform instances is* FPT *when parameterized by* nd *and $w_{\max}$, where $w_{\max} = \max\{w(e), e \in E_G\}$.*

Immediately, we get the following consequence:

Theorem 8. *For $p_1, \dots, p_k$, the $L(p_1, \dots, p_k)$-LABELING problem is* FPT *when parameterized by* nd*, k and maximum p_i (or equivalently by* nd *and the k-tuple $(p_1, \dots, p_k)$).*

Furthermore, our FPT result for CHANNEL ASSIGNMENT extends to vertex cover even without the uniformity requirement.

Theorem 9. *The* CHANNEL ASSIGNMENT *problem is* FPT *when parameterized by $w_{\max}$ and the size of vertex cover.*

One may ask whether the uniform version of CHANNEL ASSIGNMENT allows an FPT algorithm also for a broader class of graphs. Finally, we show that a natural generalization of this concept on graphs of bounded clique width yields an NP-complete problem on graphs of clique width at most 5.

2 Representing labelings as sequences and walks

We now focus on the nd-uniform instances of the CHANNEL ASSIGNMENT problem. It has been already mentioned that the optimal neighborhood diversity decomposition can be computed in cubic time. The test, whether it is nd-uniform, could be computed in quadratic additional time. On the other hand, on nd-uniform instances it suffices to consider only the type graph, whose edges take weights from the edges of the underlying graph (see Fig. 1), since such a weighted type graph corresponds uniquely to the original weighted graph, up to an isomorphism.

Hence without loss of generalization assume that our algorithms are given the type graph whose edges are weighted by separation constraints w, however we express the time complexity bounds in terms of the size of the original graph.

Without loss of generality we may assume that the given graph G and its type graph $T(G)$ are connected, since connected components can be treated independently.

If the type graph $T(G)$ contains a type t not incident with a loop, we may reduce the channel assignment problem to the graph G', obtained from G by deleting all but one vertices of the type t. Any channel assignment of G' yields a valid channel assignment of G by using the same label on all vertices of type t in G as was given to the single vertex of type t in G'. Observe that adding a loop to a type, which represents only a single vertex, does not affect the resulting graph G'. Hence we assume without loss of generality that all types are incident with a loop. We call such type graph *reflexive*.

Observation 10. *If the type graph $T(G)$ is reflexive, then vertices of G of the same type have distinct labels in every channel assignment.*

Up to an isomorphism of the graph G, any channel assignment l is uniquely characterized by a sequence of type sets as follows:

Lemma 11. *Any weighted graph G corresponding to a reflexive weighted type graph $T(G), w$ allows a channel assignment of span λ, if and only if there exists a sequence of sets $\mathcal{T} = T_0, \dots, T_\lambda$ with the following properties:*

(i) $T_i \subseteq V_{T(G)}$ *for each* $i \in [0, \lambda]$,

(ii) for each $t \in V_{T(G)} : s(t) = |\{T_i : t \in T_i\}|$,

(iii) for all $(t, r) \in E_{T(G)} : (t \in T_i \wedge r \in T_j \wedge (t \neq r \vee i \neq j)) \Rightarrow |i-j| \geq w(t, r)$

Proof. Given a channel assignment $l : V_G \to [0, \lambda]$, we define the desired sequence $\mathcal{T}$, such that the i-th element is the set of types that contain a vertex labeled by i. Formally $T_i = \{t : \exists u \in V_t : l(u) = i\}$. Now

(i) each element of the sequence is a set of types, possibly empty,

(ii) as all vertices of V_i are labeled by distinct labels by Observation 10, any type t occurs in $s(t)$ many elements of the sequence

(iii) if u of type t is labeled by i, and it is adjacent to v of type r labeled by j, then $|i - j| = |l(u) - l(v)| \geq w(u, v) = w(t, r)$, i.e. adjacent types t and r may appear in sets that are in the sequence at least $w(t, r)$ apart.

In the opposite direction assume that the sequence $\mathcal{T}$ exists. Then for each set T_i and type $t_j \in T_i$ we choose a distinct vertex $u \in V_j$ and label it by i, i.e. $l(u) = i$.

Now the condition (ii) guarantees that all vertices are labeled, while condition (iii) guarantees that all distance constraints are fulfilled. □

Observe that Lemma 11 poses no constraints on pairs of sets T_i, T_j that are at distance at least $w_{\max}$. Hence, we build an auxiliary directed graph D on all possible sequences of sets of length at most $z = w_{\max} - 1$.

The edges of D connect those sequences that overlap on a fragment of length $z-1$, i.e. when they could be consecutive in $\mathcal{T}$. This construction is well known from the so-called shift register graph.

Definition 12. *For a general graph F and weights $w : E_F \to [1, z]$ we define a directed graph D such that*

- *the vertices of V_D are all z-tuples $(T_1, \ldots, T_z)$ of subsets of V_F such that for all $(t, r) \in E_F : (t \in T_i \wedge r \in T_j) \Rightarrow |i-j| \geq w(t, r)$*
- $((T_1, \ldots, T_z), (T'_1, \ldots, T'_z)) \in E_D \Leftrightarrow T'_i = T_{i+1}$ *for all* $i \in [1, z-1]$.

As the first condition of the above definition mimics (iii) of Lemma 11 with $F = T(G)$, any sequence $\mathcal{T}$ that justifies a solution for $(T(G), w, \lambda)$, can be transformed into a walk of length $\lambda - z + 1$ in D.

In the opposite direction, namely in order to construct a walk in D, that corresponds to a valid channel assignment, we need to guarantee also an analogue of the condition (ii) of Lemma 11. In other words, each type should occur sufficiently many times in the resulting walk. Indeed, the construction of D is independent on the function s, which specifies how many vertices of each type are present in G.

In this concern we consider only special walks that allow us to count the occurrences of sets within z-tuples. Observe that V_D also contains the z-tuple $\emptyset^z = (\emptyset, \ldots, \emptyset)$. In addition, any walk of length $\lambda - z + 1$ can be converted into a closed walk from $\emptyset^z$ of length $\lambda + z + 1$, since the corresponding sequence $\mathcal{T}$ can be padded with z additional empty sets at the front, and another z empty sets at the end. From our reasoning, the following claim is immediate:

Lemma 13. *A closed walk $\mathcal{W} = W_1, \ldots, W_{\lambda+z+1}$ on D where $W_1 = W_{\lambda+z+1} = \emptyset^z$, yields a solution of the* Channel Assignment *problem on a* nd-*uniform instance G, w, λ with reflexive $T(G)$, if and only if $s(t) = |\{W_i : t \in (W_i)_1\}|$ holds for each $t \in V_{T(G)}$.*

We found interesting that our representation of the solution resembles the NP-hardness reduction found by Griggs and Yeh [13] (it was briefly outlined in Section 1.4) and later generalized by Bodlaender et al. [1]. The key difference is that in their reduction, a Hamilton path is represented by a sequence of vertices of the constructed graph. In contrast, we consider walks in the type graph, which is assumed to be of limited size.

3 The algorithm

In this section we prove the following statement, which directly implies our main result, Theorem 7:

Proposition 14. *Let G, w be a weighted graph, whose weights are uniform with respect to a neighborhood diversity partition with τ classes.*

Then the Channel Assignment *problem can be decided on G, w and any λ in time $2^{2^{O(\tau w_{\max})}} \log n$, where n is the number of vertices of G, provided that G, w are described by a weighted type graph $T(G)$ on τ nodes.*

A suitable labeling of G can be found in additional $2^{2^{O(\tau w_{\max})}} n$ time.

Proof. According to Lemma 13, it suffices to find a closed walk $\mathcal{W}$ (if it exists) corresponding to the desired labeling l. From the well-known Euler's theorem it follows that any directed closed walk $\mathcal{W}$ yields a multiset of edges in D that induces a connected subgraph and that satisfies Kirchhoff's law. In addition, any such suitable multiset of edges can be converted into a closed walk, though the result need not be unique.

For this purpose we introduce an integer variable $\alpha_{(W,U)}$ for every directed edge $(W, U) \in E_D$. The value of the variable $\alpha_{(W,U)}$ is the number of occurrences of (W, U) in the multiset of edges.

Kirchhoff's law is straightforwardly expressed as:

$$\forall W \in V_D : \sum_{U:(W,U)\in E_D} \alpha_{(W,U)} - \sum_{U:(U,W)\in E_D} \alpha_{(U,W)} = 0$$

In order to guarantee connectivity, observe first that an edge (W, U) and $\emptyset^z$ would be in distinct components of a subgraph of D, if the subgraph is formed by removing edges that include a cut C between (W, U) and $\emptyset^z$. Now, the chosen multiset of edges is disconnected from $\emptyset^z$, if there is such an edge (W, U) together with a cut set C such that $\alpha_{(W,U)}$ has a positive value, while all variables corresponding to elements of C are zeros. As all variable values are bounded above by λ, we express that C is not a cutset for the chosen multiset of edges by the following condition:

$$\alpha_{(W,U)} - \lambda \sum_{e\in C} \alpha_e \leq 0$$

To guarantee the overall connectivity, we apply the above condition for every edge $(W, U) \in E_D$, where $W, U \neq \emptyset^z$, and for each set of edges C that separates W or U from $\emptyset^z$.

The necessary condition expressed in Lemma 13 can be stated in terms of variables $\alpha_{(W,U)}$ as

$$\forall t \in V_{T(G)} : \sum_{W: t \in (W)_1} \sum_{U:(W,U) \in E_D} \alpha_{(W,U)} = s(t)$$

Finally, the size of the multiset is the length of the walk, i.e.

$$\sum_{(W,U) \in E_D} \alpha_{(W,U)} = \lambda + z + 1$$

Observe that these conditions for all (W,U) and all suitable C indeed imply that the $\emptyset^z$ belongs to the subgraph induced by edges with positively evaluated variables $\alpha_{(W,U)}$.

Algorithm 1 summarizes our deductions.

To complete the proof, we argue about the time complexity as follows:

- Line 1 needs $O(|E_{T(G)}|) = O(\tau^2)$ time.
- As D has at most $2^{\tau z}$ nodes and at most $2^{\tau(z+1)}$ edges, line 2 needs $2^{O(\tau z)}$ time.
- Similarly, conditions at lines 4 and 5 require $2^{O(\tau z)}$ time and space to be composed. Analogously, conditions at lines 7 and 8 involve coefficients that are proportional to the size of the original graph G (namely λ and $s(t)$), hence $2^{O(\tau z)} \log n$ time and space is needed here.
- For line 6, we examine each subset of E_D and decide whether it is a suitable cutset C. There are at most $2^{2^{\tau(z+1)}}$ choices for C, so the overall time and space complexity for the composition of conditions at line 6 is $2^{2^{O(\tau z)}} \log n$.
- Frank and Tardos [10] (improving the former result due to Lenstra [18]) showed that the time needed to solve the system of inequalities with p integer variables is $O(p^{2.5p+o(p)}L)$, where L is the number of bits needed to encode the input. As we have $2^{O(\tau z)}$ variables and the conditions are encoded in space $2^{2^{O(\tau z)}} \log n$, the time needed to resolve the system of inequalities is $2^{2^{O(\tau z)}} \log n$.
- A solution of the ILP can be converted into the walk in time $2^{2^{O(\tau z)}} n$, and the same bound applies to the conversion of a walk to the labeling at lines 10 and 11.

Input: A reflexive type graph $T(G)$ whose edges are labeled by w and a span λ.
Output: A channel assignment $l : G \to [0, \lambda]$ respecting constraints w, if it exists.

begin

1 Compute $z := w_{\max} - 1$;

2 Construct the directed graph D;

3 Solve the following ILP in variables $\alpha_{(W,U)} : (W, U) \in E_D$:

4 for each $(W, U) \in E_D$:

$$\alpha_{(W,U)} \geq 0$$

5 for each $W \in V_D$:

$$\sum_{U:(W,U)\in E_D} \alpha_{(W,U)} - \sum_{U:(U,W)\in E_D} \alpha_{(U,W)} = 0$$

6 for each $(W, U) \in E_D$ and each cutset C between (W, U) and $\emptyset^z$ in D:

$$\alpha_{(W,U)} - \lambda \sum_{e\in C} \alpha_e \leq 0$$

7 for each $t \in V_{T(G)}$:

$$\sum_{W:t\in(W)_1} \sum_{U:(W,U)\in E_D} \alpha_{(W,U)} = s(t)$$

8 $$\sum_{(W,U)\in E_D} \alpha_{(W,U)} = \lambda + z + 1;$$

9 **if** *the ILP has a solution* **then**

10 find a walk $\mathcal{W}$ that traverses each edge (W, U) exactly $\alpha_{(W,U)}$ times;

11 convert the walk $\mathcal{W}$ into a labeling l and **return** l;

else

12 **return** *"No channel assignment l of span λ exists."*

end

end

Algorithm 1: Solving the Channel Assignment problem.

Observe that if only the existence of the labeling should be decided, the lines 10 and 11 need not to be executed, only an affirmative answer needs to be returned instead. □

We are aware the the double exponential dependency on nd and $w_{\max}$

makes our algorithm interesting mostly from the theoretical perspective. Naturally, one may ask, whether the exponential tower height might be reduced or whether some nontrivial lower bounds on the computational complexity could be established (under usual assumptions on classes in the complexity hierarchy).

4 Bounded vertex-cover

We utilize the results of the previous sections to derive an FPT algorithm proposed as Theorem 9.

Proof of Theorem 9. Given a graph G and its optimal vertex cover U, we construct a partition of the vertices of G as follows. Next let $I = V(G) \setminus U$ be the independent set of G. We form a partition of I as follows: For every subset $X \subseteq U$ we define:

$$I_X := \{v \in I : \{v, x\} \in E(G) \text{ for } x \in X \text{ and } \{v, x\} \notin E(G) \text{ for } x \notin X\}.$$

Observe that for any $u, v \in I_X$ it holds that $N(u) = X = N(v)$, and hence also $u \sim v$. In particular, the optimal neighborhood diversity decomposition of G consists of all nonempty sets I_X together with a suitable partition of U. Consequently, $\mathrm{nd}(G) \leq 2^{\mathrm{vc}(G)} + \mathrm{vc}(G)$.

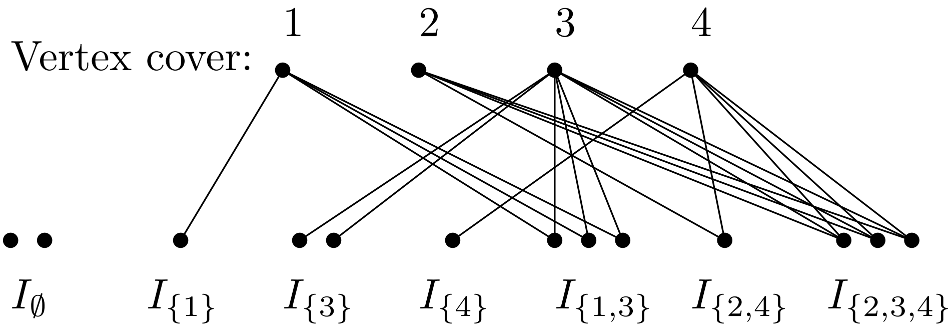

Figure 2: An example of a neighborhood diversity decomposition based on a vertex cover.

We further refine sets I_X, so that the edge-weights became uniform. For a set $X = \{x_1, \ldots, x_k\} \subseteq U$ and each k-tuple of positive integers $\mathbf{w} = (w_1, \ldots, w_k)$ with $0 < w_i \leq w_{\max}$ for every $1 \leq i \leq k$ we define the set $I_X^{\mathbf{w}}$ as

$$I_X^{\mathbf{w}} := \{v \in I_X : w(v, x_i) = w_i \text{ for } 1 \leq i \leq k\}.$$

Observe that any refinement of a neighborhood diversity decomposition is again a decomposition. We now estimate the number of types of the refined decomposition. The number of types of the refined decomposition can be upper-bounded by $\mathrm{vc}(G) + (2^{\mathrm{vc}(G)}) w_{\max}^{\mathrm{vc}(G)}$. To finish the proof we apply Proposition 14 on the refined decomposition. □

5 NLC-uniform channel assignment

One may ask whether the concept of nd-uniform weights could be extended to broader graph classes. We show, that already its direct extension to graphs of bounded clique width makes the Channel Assignment problem NP-complete. Instead of clique width we express our results in terms of NLC-width [21] (NLC stands for node label controlled). The parameter NLC-width is linearly dependent on clique width, but it is technically simpler.

We now briefly review the related terminology. A NLC-decomposition of a graph G is a rooted tree whose leaves are in one-to-one correspondence with the vertices of G. For the purpose of inserting edges, each vertex is given a label (the labels for channel assignment are now irrelevant), which may change during the construction of the graph G. Internal nodes of the tree are of two kinds: *relabel* nodes and *join* nodes.

Each relabel node has a single child and as a parameter takes a mapping ρ on the set of labels. The graph corresponding to a relabel node is isomorphic to the graph corresponding to its child, only ρ is applied on each vertex label.

Each join node has a two children and as a parameter takes a binary relation S on the set of labels. The graph corresponding to a join node is isomorphic to the disjoint union of the two graphs G_1 and G_2 corresponding to its children, where further edges are inserted as follows: $u \in V_{G_1}$ labeled by i is made adjacent to $v \in V_{G_2}$ labeled by j if and only if $(i, j) \in S$.

The minimum number of labels needed to construct at least one labeling of G in this way is the NLC width of G, denoted by $\mathrm{nlc}(G)$.

Observe that $\mathrm{nlc}(G) \leq \mathrm{nd}(G)$ as the vertex types could be used as labels for the corresponding vertices and the adjacency relation in the type graph could be used for S in all join nodes. In particular, in this construction the order of performing joins is irrelevant and no relabel nodes are needed.

Definition 15. *The edge weights w on a graph G are* nlc-uniform *with respect to a particular NLC-decomposition, if $w(u, v) = w(u', v')$ whenever edges (u, v) and (u', v') are inserted during the same join operation and at the* *moment of insertion u, u' have the same label in G_1 and v, v' have the same label in G_2.*

Observe that our comment before the last definition justifies that weights that are uniform with respect to a neighborhood diversity decomposition are uniform also with respect to the corresponding NLC-decomposition.

Gurski and Wanke showed that the NLC-width remains bounded when taking powers of trees [14]. It is well known that NLC-width of a tree is at most three. Fiala et al. proved that $L(3,2)$-LABELING is NP-complete on trees [6]. To combine these facts together we show that the weights on the graph arising from a reduction of the $L(3,2)$-labeling on a tree to CHANNEL ASSIGNMENT are nlc-uniform.

Theorem 16. *The* CHANNEL ASSIGNMENT *problem is* NP*-complete on graphs with edge weights that are* nlc*-uniform with respect to an NLC-decomposition of width at most four.*

Proof. Let a tree T be an instance of the $L(3,2)$-LABELING problem.

By induction on the size of T we show that T^2 allows an NLC-decomposition such that the weights w prescribed by the reduction of $L(3,2)$-LABELING to CHANNEL ASSIGNMENT are nlc-uniform.

Assume for the induction hypothesis that such an NLC-decomposition exists for every tree T' on less than n vertices, where the labels of T' are distributed as follows: assume that T' is rooted in a vertex r', then r' is labeled by 1, its direct neighbors by 2 and all other vertices by 3.

Such a decomposition clearly exists for a tree on a single vertex.

Now consider a tree T on $n \geq 2$ nodes. Choose an edge (r', r'') arbitrarily and define two trees T' and T'' as the components of $T \setminus (r', r'')$, where T' contains r' and vice versa.

By the induction hypothesis T' and T'' allow NLC-decompositions with the desired properties. Before we join trees T' and T'' together, we change labels in T'' as $1 \to 2, 2 \to 4$. At the join we will insert the following weighted edges: of weight 3 between vertices labeled 1 in T' and 2 in T'', and of weight 2 between vertices of labels 2 and 2, and between 1 and 4, respectively. Finally, we relabel $4 \to 3$ and promote r' to be the root r of T. All steps are depicted in Fig. 3. Observe that the result of this construction is T^2 with appropriate nlc-uniform weights, and that its labeling satisfies all conditions of the induction hypothesis. □

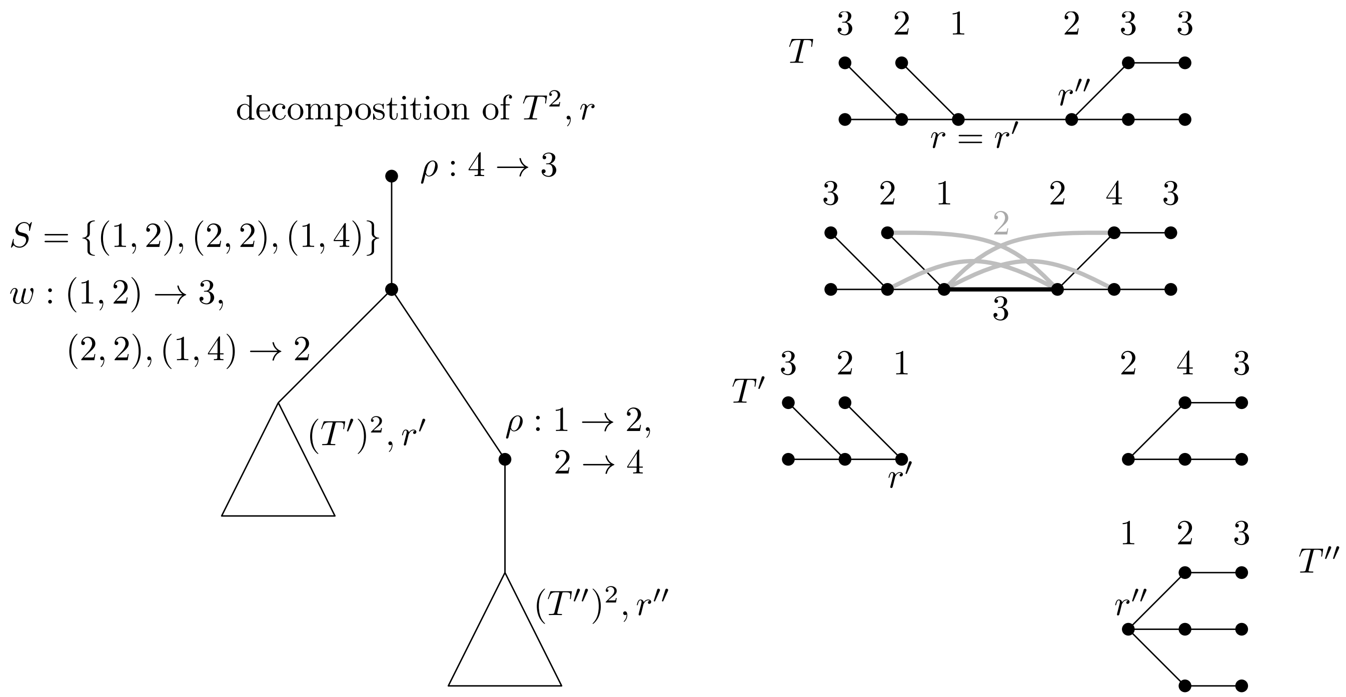

Figure 3: Recursive step in the construction of an NLC-decomposition. The original edges of T, T' and T'' are in black. The weighted edges added in this step are in bold.

6 Conclusion

We have shown an algorithm for the Channel Assignment problem on nd-uniform instances and several complexity consequences for the $L(p_1, \ldots p_k)$-labeling problem. In particular, Theorem 8 extends known results for the $L(p,1)$-labeling problem to labelings with arbitrarily many distance constraints, answering an open question of [7]. Simultaneously, we broaden the considered graph classes by restricting neighborhood diversity instead of vertex cover.

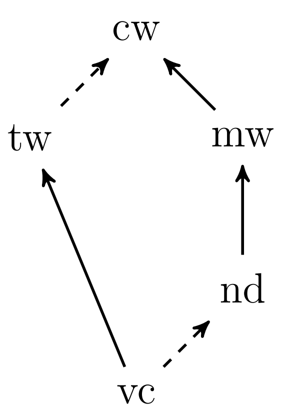

Figure 4: A map of assumed parameters. Full arrow stands for linear upper bounds, while dashed arrow stands for exponential upper bounds.

While the main technical tools of our algorithms are bounded-dimension

ILP programs, ubiquitous in the FPT area, the paper shows an interesting insight on the nature of the labelings over the type graph and the necessary patterns of such labelings of very high span. Note that the span of a graph is generally not bounded by any of the considered parameters and may be even proportional to the order of the graph.

Solving a generalized problem on graphs of bounded neighborhood diversity is a viable method for designing FPT algorithms for a given problem on graphs of bounded vertex cover, as demonstrated by this and previous papers. This promotes neighborhood diversity as a parameter that naturally generalizes the widely studied parameter vertex cover.

We would like to point out that the parameter *modular width*, proposed by Gajarský, Lampis and Ordyniak [11], offers further generalization of neighborhood diversity towards the clique width [4] (dependencies between these graph parameters are depicted in Fig. 4).

As an interesting open problem we ask whether it is possible to strengthen our results to graphs of bounded modular width or whether the problem might be already NP-complete for fixed modular width, as is the case with clique width. For example, the Graph Coloring problem ILP based algorithm for bounded neighborhood diversity translates naturally to an algorithm for bounded modular width. On the other hand, there is no apparent way how our labeling results could be adapted to modular width in a similar way.

We thank Andrzej Proskurowski and Tomáš Masařík for valuable comments.

References

[1] Hans L. Bodlaender, Ton Kloks, Richard B. Tan, and Jan van Leeuwen. λ-coloring of graphs. In Horst Reichel and Sophie Tison, editors, *STACS*, volume 1770 of *Lecture Notes in Computer Science*, pages 395–406. Springer, 2000.

[2] Gerald J. Chang and David Kuo. The $L(2,1)$-labeling problem on graphs. *SIAM Journal on Discrete Mathematics*, 9(2):309–316, 1996.

[3] Derek C. Corneil and Udi Rotics. On the relationship between clique-width and treewidth. *SIAM J. Comput.*, 34(4):825–847, 2005.

[4] Bruno Courcelle and Stephan Olariu. Upper bounds to the clique width of graphs. *Discrete Appl. Math.*, 101(1–3):77–114, 2000.

[5] Jiří Fiala, Petr A. Golovach, and Jan Kratochvíl. Distance constrained labelings of graphs of bounded treewidth. In Luís Caires, Giuseppe F. Italiano, Luís Monteiro, Catuscia Palamidessi, and Moti Yung, editors, *ICALP*, volume 3580 of *Lecture Notes in Computer Science*, pages 360–372. Springer, 2005.

[6] Jiří Fiala, Petr A. Golovach, and Jan Kratochvíl. Computational complexity of the distance constrained labeling problem for trees (extended abstract). In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *ICALP (1)*, volume 5125 of *Lecture Notes in Computer Science*, pages 294–305. Springer, 2008.

[7] Jiří Fiala, Petr A. Golovach, and Jan Kratochvíl. Parameterized complexity of coloring problems: Treewidth versus vertex cover. *Theor. Comput. Sci.*, 412(23):2513–2523, 2011.

[8] Jiří Fiala, Jan Kratochvíl, and Ton Kloks. Fixed-parameter complexity of λ-labelings. *Discrete Applied Mathematics*, 113(1):59–72, 2001.

[9] Fedor V. Fomin, Petr A. Golovach, Daniel Lokshtanov, and Saket Saurabh. Intractability of clique-width parameterizations. *SIAM Journal of Computing*, 39(5):1941–1956, 2010.

[10] András Frank and Éva Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7(1):49–65, 1987.

[11] Jakub Gajarský, Michael Lampis, and Sebastian Ordyniak. Parameterized algorithms for modular-width. In Gregory Gutin and Stefan Szeider, editors, *Parameterized and Exact Computation - 8th International Symposium, IPEC 2013, Sophia Antipolis, France, September 4-6, 2013, Revised Selected Papers*, volume 8246 of *Lecture Notes in Computer Science*, pages 163–176. Springer, 2013.

[12] Robert Ganian and Jan Obdržálek. Expanding the expressive power of monadic second-order logic on restricted graph classes. In Thierry

Lecroq and Laurent Mouchard, editors, *Combinatorial Algorithms - 24th International Workshop, IWOCA 2013, Rouen, France, July 10-12, 2013, Revised Selected Papers*, volume 8288 of *Lecture Notes in Computer Science*, pages 164–177. Springer, 2013.

[13] Jerrold R. Griggs and Roger K. Yeh. Labelling graphs with a condition at distance 2. *SIAM Journal on Discrete Mathematics*, 5(4):586–595, 1992.

[14] Frank Gurski and Egon Wanke. The NLC-width and clique-width for powers of graphs of bounded tree-width. *Discrete Applied Mathematics*, 157(4):583–595, 2009.

[15] Daniel Kobler and Udi Rotics. Polynomial algorithms for partitioning problems on graphs with fixed clique-width (extended abstract). In *Symposium on Discrete algorithms, 12th SODA'01, Washington*, pages 468–476. ACM-SIAM, 2001.

[16] Daniel Král. The channel assignment problem with variable weights. *SIAM J. Discrete Math.*, 20(3):690–704, 2006.

[17] Michael Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012.

[18] H. W. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.

[19] Colin McDiarmid and Bruce Reed. Channel assignment on graphs of bounded treewidth. *Discrete Mathematics*, 273(1–3):183–192, 2003.

[20] Marián Škvarek. The channel assignment problem for series-parallel graphs. Bachelor's thesis, Charles University, Prague, 2010. in Czech.

[21] Egon Wanke. k-NLC graphs and polynomial algorithms. *Discrete Applied Mathematics*, 54(2–3):251–266, 1994.