Distance constrained labelings of graphs of bounded treewidth *

Jiří Fiala[a], Petr Golovach[b], Jan Kratochvíl[a]

[a] Institute for Theoretical Computer Science[1]
and Department of Applied Mathematics,
Faculty of Mathematics and Physics,
Charles University, Prague, Czech Republic
{fiala,honza}@kam.mff.cuni.cz

[b] Matematicheskii Fakultet,
Syktyvkar State University,
Syktyvkar, Russia
golovach@ssu.komi.com

Abstract

We prove that the L(2,1)-LABELING problem is NP-complete for graphs of treewidth two, thus adding a natural and well studied problem to the short list of problems whose computational complexity separates treewidth one from treewidth two. We prove similar results for other variants of the distance constrained graph labeling problem.

1 Introduction

The notion of distance constrained graph labeling attracted a lot of attention in the past years both for its motivation by the practical frequency assignment problem, and for its interesting graph theoretic properties. The task of

*Research supported in part by project KONTAKT 525 — DIMACS-DIMATIA-Rényi Cooperation in Discrete Mathematics

[1]Supported by the Ministry of Education of the Czech Republic as project 1M0021620808.

assigning frequencies to transmitters to avoid undesired interference of signals is modeled in several ways. The so called *channel assignment problem* assumes that a minimum allowed difference of channels is given for every two transmitters. Thus the input of this problem is a weighted graph whose vertices correspond to the transmitters, and the task is to assign nonnegative integers (channels) to the vertices so that for every edge, the difference of the assigned channels is at least the weight of the edge, and so that the largest channel used is minimized.

Another approach, and this one we follow in the present paper, is the *distance constrained graph labeling.* Here it is assumed that the distance of transmitters can be modeled by a graph, and that the distance of the transmitters influences possible interference in such a way that the closer two transmitters are, the farther apart their frequencies must be. Formally, an assignment of nonnegative integers to the vertices of a graph G is an $L(p_1, \ldots, p_k)$-labeling if for every two vertices at distance at most $i \leq k$, the difference of the integers (labels) assigned to them is at least p_i. Here $k \geq 1$ is the depth to which the distance constraints are applied, and integers $p_1 \geq p_2 \geq \ldots \geq p_k$ are parameters of the problem. Again, the goal is to minimize the maximum label used. The most studied of the distance constrained labelings is the case $k = 2, p_1 = 2, p_2 = 1$, i.e., the $L(2,1)$-*labeling*. In this case adjacent vertices must be assigned labels that differ by at least 2, while nonadjacent vertices with a common neighbor must be assigned distinct labels. The maximum label used is called the *span* of the labeling. The minimum span of an $L(2,1)$-labeling of a graph G will be denoted by $L_{(2,1)}(G)$.

The notion of $L(2,1)$-labeling was in fact first proposed by Roberts [20] and many nontrivial results were presented in a pioneer paper of Griggs and Yeh [15]. Let us mention their conjecture that $L_{(2,1)}(G) \leq \Delta^2(G)$ (where $\Delta(G)$ stands for the maximum vertex degree in G). This conjecture has been verified for various graph classes, but it is still open for general graphs (with $L_{(2,1)}(G) \leq \Delta(G)^2 + \Delta(G) - 1$ being the current record [16]). From the computational complexity point of view, Griggs and Yeh proved that determining $L_{(2,1)}(G)$ is an NP-hard problem, and this result was later strengthened by Fiala et al. [7] by showing that deciding $L_{(2,1)}(G) \leq k$ is NP-complete for every fixed $k \geq 4$. Griggs and Yeh also conjectured that it is NP-complete to compute the $L_{(2,1)}$ number of a tree, but this was somewhat surprisingly disproved by a dynamic programming polynomial time algorithm of Chang and Kuo [4].

The common expectation says that problems solvable in polynomial time for trees should also be polynomially solvable for graphs of bounded treewidth, though sometimes the extension to bounded treewidth is not straightforward (cf. e.g., the case of chromatic index [2]). (We informally recall that the treewidth is a graph invariant that describes how far is the graph from being a tree. For a formal definition the reader is referred to a survey [3] or to one of the original papers [1] introducing this invariant in terms of so called *partial k-trees*. For our purposes we only need the fact that graphs of treewidth at most two are exactly the graphs that do not contain a subdvision of K_4 as a subgraph, and connected graphs of treewidth one are exactly trees.) Only very few exceptions to this rule of thumb are known, and in fact very few problems are known to be hard for graphs of bounded treewidth. An example is, e.g., the MINIMUM BANDWIDTH problem (which is NP-hard already for trees [12]) or the closely related CHANNEL ASSIGNMENT problem which has been recently shown NP-complete for graphs of treewidth three [18]. The natural question of the complexity of $L(2,1)$-labelings for graphs of bounded treewidth has been posed many times and remained open since 1996. The main result of our paper settles it by showing that determining the $L_{(2,1)}$ number of graphs of treewidth two is NP-hard.

Before we formulate the result formally, we specify precisely what problem we deal with. The decision problem whether a given graph admits an $L(2,1)$-labeling of *fixed* span can be described in Monadic Second Order Logic (MSOL), and therefore is solvable in linear time for any class of graphs of bounded treewidth by a generic algorithm of Courcelle [5]. Thus we naturally assume that the span is a part of the input, and we consider the following problem.

$L(2,1)$-LABELING

Input: An integer λ and a graph G.

Question: Is $L_{(2,1)}(G) \leq \lambda$?

Theorem 1. *The* L(2,1)-LABELING *problem is* NP*-complete for graphs of treewidth at most two.*

So far we have only discussed the model in which interference of the frequencies (or channels) decreases linearly with their increasing difference. It is, however, plausible to consider also such models in which frequencies far apart may interfere (e.g., if one is a multiple of the other one). This means more complicated metrics in the frequency space. A concrete step in this direction is the cyclic metric introduced by van Heuvel et al. [23]. In this metric, the graph of the channel space is the cycle of length λ. Similarly to the linear case, we talk about $C(2,1)$-labelings and denote by $C_{(2,1)}(G)$ the minimum span of a $C(2,1)$-labeling of G (note that in the cyclic metric, the span is the number of available channels, not the difference between the largest and smallest one). For general graphs, deciding if $C_{(2,1)}(G) \leq \lambda$ is NP-complete for every fixed $\lambda \geq 6$ [9]. For λ part of the input and graphs of bounded treewidth, we fully characterize the complexity of the C(2,1)-LABELING problem (which, given a graph G and an integer λ as input, asks if $C_{(2,1)}(G) \leq \lambda$):

Proposition 2. **([17])** *Let T be a tree with at least one edge, and $p \geq q$ nonnegative integers. Then*

$$C_{(p,q)}(T) = q\Delta(T) + 2p - q$$

where $\Delta(T)$ is the maximum degree of a vertex in T.

Theorem 3. *The* C(2,1)-LABELING *problem is* NP*-complete for graphs of treewidth at most two.*

Fiala and Kratochvíl [9] defined the notion of $H(2,1)$-labeling as the utmost generalization in the case when the metric of the channel space can be described by a graph H, and showed that $H(2,1)$-labelings of a graph G are exactly locally injective homomorphisms from G into the complement of H. The complexity of the H(2,1)-LABELING problem for some parameter graphs H then follows from [8], but the complete characterization is not even in sight. On the other hand, if G has bounded treewidth, the H(2,1)-LABELING problem is solvable in polynomial time since for a fixed graph H, the existence of an $H(2,1)$-labeling of G can be expressed in MSOL.

It remains to study thecase when both G and H are part of the input and we refer to it as the(2,1)-LABELING problem. Observe that the L(2,1)-LABELING problem is the restriction of (2,1)-LABELING to inputs such that H is a path. Hence it followsfrom Theorem 1 that (2,1)-LABELING is NP-complete for graphs of treewidth two. However, in this most general setting, we are able to prove dichotomy even with respect to pathwidth (for definition of pathwidth see [21, 22, 3], just recall that connected graphs of pathwidth one are exactly caterpillars):

Theorem 4. *For a tree T with m vertices and an arbitrary graph H with n vertices, one can decide in time $O(n^3m^2)$ whether T allows an $H(2,1)$-labeling.*

Theorem 5. *The* (2,1)-LABELING *problem is* NP*-complete for graphs G of pathwidth at most two (the graph H may be arbitrary).*

The paper is organized as follows. In Section 2 we review technical definitions and notation and prove an auxiliary result on systems of distant representatives for symmetric sets. The main result, Theorem 1, is proved in Section 3. The case of cyclic metric is discussed in Section 4. Theorems 4 and 5 are proved in Section 5. The last section contains concluding remarks and open questions.

2 Preliminaries

All graphs considered are finite and simple, i.e., with a finite vertex set and without loops or multiple edges. For a vertex u, the symbol $N(u)$ denotes the *open neighborhood* of u, i.e., the set of all vertices adjacent to u, and we denote by $\deg u = |N(u)|$ the *degree* of u.

A graph is called *series-parallel* if it can built from isolated edges with endvertices called South and North poles by a sequence of series and parallel compositions (the former identifies the North pole of one component with the South pole of the other one, the latter unifies the North poles of the components into a common North pole, and likewise the South poles). It is well known that a graph has treewidth at most two if and only if all its 2-connected subgraphs are series-parallel.

The labels are always nonnegative integers, with 0 being the smallest label used. We use the notation $[x, y] = \{x, x + 1, \ldots, y - 1, y\}$ to denote intervals of consecutive integers. We say that a set S of integers is *symmetric within* an interval $[x, y]$ if $S \subseteq [x, y]$ and for every $i \in [x, y]$, $i \in S$ if and only if $y + x - i \in S$.

A system of distinct representatives for a set system $S_1, S_2, \ldots, S_n$ is a system of distinct elements $s_i \in S_i, i = 1, 2, \ldots, n$. The theory of SDR's is well developed, the necessary and sufficient condition for their existence is given by the well known Hall theorem, and an SDR can be found in polynomial time (e.g., by a bipartite matching algorithm). If the ground set $\bigcup_{i=1}^{n} S_i$ is equipped with a metric function, we can further impose conditions on the distance of the chosen representatives. We refer the reader to [11, 14] for a survey on the computational complexity of finding systems of distant representatives for sets in metric spaces and their applications in various graph labeling problems. Now we will use a special variant of this problem as an auxiliary tool:

SRL (Special representatives in the linear metric)

Input: An integer n and a collection of sets of integers $S_1, S_2, \ldots, S_m$ symmetric within the interval $[2, \lambda - 2]$, where $\lambda = 4n + 5$.

Question: Does there exist a collection of distinct integers $s_1, s_2, \ldots, s_m, t_1, t_2, \ldots, t_n, u_1, u_2, \ldots, u_n$ such that

- $s_i \in S_i$ for every $i = 1, \ldots, m$,
- $t_i \in \{2i, \lambda - 2i - 1\}$ for every $i = 1, \ldots, n$,
- $u_i \in \{2i + 1, \lambda - 2i\}$ for every $i = 1, \ldots, n$,
- $|t_i - u_i| \geq 2$ for every $i = 1, 2, \ldots, n$?

Lemma 2.1. *The problem* SRL *is* NP*-complete.*

The proof is based on the following special variant of the 3-SAT problem (known NP-complete, cf. e.g. [6]).

2-3-SAT

Input: A Boolean formula Φ in conjunctive normal form, whose each clause consists of 2 or 3 literals and whose every variable has at most 2 positive and at most 2 negative occurrences.

Question: Is Φ satisfiable?

Proof. We reduce from 2-3-SAT (cf. Section 2). Let Φ have n variables $x_1, \ldots, x_n$ and m clauses $C_1, \ldots, C_m$. The number of variables n will be the n from the input of SRL. Recall that $\lambda = 4n + 5$. For every $j = 1, 2, \ldots, m$ the set S_j is constructed from the clause C_j as follows

$$S_j = \bigcup_{i: x_i \in C_j} \{2i, \lambda - 2i\} \cup \bigcup_{i: \neg x_i \in C_j} \{2i + 1, \lambda - 2i - 1\}.$$

Thus every set S_j has 4 or 6 elements and is symmetric within $[0, \lambda]$.

Assume that Φ allows a satisfying assignment. If a variable x_i is assigned the value true, we set $t_i = 2i - 1$, $u_i = \lambda - 2i - 1$. Analogously for x_i negatively valued, we let $t_i = 2i$, $u_i = \lambda - 2i$. For each clause C_j we choose one satisfying literal. If C_j is satisfied by the literal x_i for some $i = 1, 2, \ldots, n$, we let $s_j = 2i$, if x_i is the first occurrence of x_i in Φ, and $s_j = \lambda - 2i$ for the second occurrence of x_i in Φ. In the case C_j is satisfied by $\neg x_i$ we choose $s_j = 2i + 1$ for the first occurrence of $\neg x_i$ and $s_j = \lambda - 2i - 1$ otherwise. Straightforwardly, the collection $s_1, \ldots, u_n$ satisfies all four properties from the definition of the SRL problem.

For the opposite direction suppose that $s_1, \ldots, u_n$ is a valid solution for the SRL problem. The crucial observation is that for every $i = 1, 2, \ldots, n$, there are only two possible choices for the values of t_i and u_i so that $|t_i - u_i| \geq 2$. Namely, either $t_i = 2i$ and $u_i = \lambda - 2i$ or alternatively $t_i = 2i + 1$ and $u_i = \lambda - 2i - 1$. In the first case we assign $x_i =$ false and accordingly $x_i =$ true in the second case.

Then for each $j = 1, \ldots, m$, the value of u_j indicates the satisfying literal for the clause C_j: If $u_j = 2i$ or $\lambda - 2i$, then C_j is satisfied by the true assignment to the variable x_i. Alternatively, if $u_j = 2i + 1$ or $\lambda - 2i - i$ then the literal $\neg x_i$ satisfies C_j as the variable x_i is assigned false.

Since the size of the family $S_1, S_2, \ldots, S_m$ is polynomial in the size of Φ, 2-3-SAT$\propto$SRL as claimed. $\square$

We continue with an analogous lemma for set systems over a space with the cyclic metric: For a fixed $\lambda \geq 2$ and integers $x, y \in \{0, 1, \ldots, \lambda - 1\}$ we denote by $\rho_c(x, y) = \min\{|x - y|, \lambda - |x - y|\}$ the *(cyclic) distance* between x and y.

SRC (Special representatives in the cyclic metric)

Input: An integer n and a collection of sets of integers $S_1, S_2, \ldots, S_m$ symmetric within the interval $[1, \lambda - 1]$ for $\lambda = 16n$.

Question: Does there exists a collection of distinct integers $s_1, s_2, \ldots, s_m, t_1, t_2, \ldots, t_n, u_1, u_2, \ldots, u_n$ such that

- for all $i = 1, \ldots, m : s_i \in S_i$,
- for all $i = 1, \ldots, n : t_i \in \{4n - i, 4n + i\}$,
- for all $i = 1, \ldots, n : u_i \in \{12n - i, 12n + i\}$,

- $\rho_c(t_i, u_i) = 8n - 2i$ for every $i = 1, 2, \ldots, n$?

Lemma 2.2. *The problem* SRC *is* NP*-complete.*

Proof. We show 2-3-SAT $\propto$ SRC.

For every $j = 1, 2, \ldots, m$ the set S_j is constructed from the clause $C_j \in \Phi$ by following rules: if the clause C_j contains the literal x_i then integers $4n - i, 12n + i$ are inserted into S_j, and if $\neg x_i \in C_j$ then we put integers $4n + i, 12n - i$ into S_j. It can easily be seen that each set S_j is symmetric.

Assume that Φ allows a satisfying assignment. We choose $t_i = 4n + i$ and $u_i = 12n - i$ when a variable x_i is assigned true and we let $t_i = 4n - i$ and $u_i = 12n + i$ otherwise. From each clause C_j we pick one satisfying literal. When C_j is satisfied by x_i we select $s_i = 4n - i$ if x_i is the first occurrence of x_i in Φ, and $s_i = 12n + i$ if it is the second occurrence. In the case $\neg x_i$ satisfies C_j we choose $s_i = 4n + i$ for the first occurrence of $\neg x_i$ in Φ and $c_i = 12n - i$ for the second.

It can be easily seen that all $s_1, \ldots, u_n$ are distinct and $\rho_c(t_i - u_i) = 8n - 2i$ for every $i = 1, 2, \ldots, n$. Hence, it is the desired solution of the SRC problem.

Now suppose $s_1, \ldots, u_n$ satisfy all four conditions of the SRC problem. As in the prof of the previous lemma, there are only two possible choices for pairs (t_i, u_i): as $\rho_c(t_i, u_i) = 8n - 2i > 2i$ either $(t_i, u_i) = (4n - i, 12n + i)$ (indicating x_i being assigned false) or $(t_i, u_i) = (4n + i, 12n - i)$ (x_i being evaluated true). Consequently, for each $j = 1, \ldots, m$ the choice of u_j provides a satisfying literal for the clause C_j: If $u_j = 4n - i$ or $12n + i$ then C_j is satisfied as x_i is assigned true. Analogously, when $u_j = 4n + i$ or $12n - i$ then the literal $\neg x_i$ satisfies the clause C_j. □

3 $L(2, 1)$-labeling of graphs of treewidth two

This entire section is devoted to the proof of Theorem 1. We will utilize Lemma 2.1 and reduce from the SRL problem. Suppose we are given integers n and $\lambda = 4n + 5$, and m subsets $S_1, \ldots, S_m$ of $[2, \lambda - 2]$ which are all symmetric within this interval (we may further assume that all of them have size at most 6, but this is not important for our proof). Our aim is to construct a graph G' of treewidth two such that $L_{(2,1)}(G') \leq \lambda$ if and only if the given instance of SRL is feasible. The construction of G' is achieved in several steps.

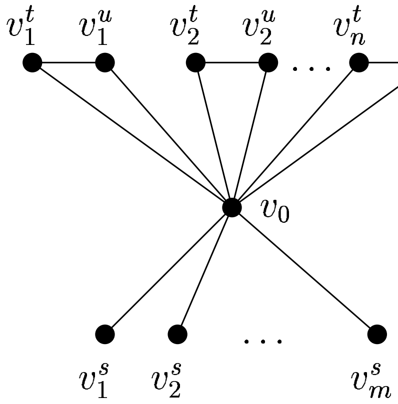

Figure 1: The graph G.

3.1 Reduction to List Labeling Construct the graph G on vertices $V_G = (v_0, v_1^s, \ldots, v_m^s, v_1^t, \ldots, v_n^t, v_1^u, \ldots, v_n^u)$ where v_0 is adjacent to all other vertices, and futhermore $(v_i^t, v_i^u) \in E_G$ for all $i = 1, \ldots, n$. (See Fig. 1.) To each vertex of $x \in V_G$ we assign a set of admissible labels as follows

- $T(v_0) = \{0, \lambda\}$
- $T(v_i^s) = S_i$ for all $i = 1, \ldots, m$
- $T(v_i^t) = \{2i, \lambda - 2i - 1\}$ for all $i = 1, \ldots, n$
- $T(v_i^u) = \{2i + 1, \lambda - 2i\}$ for all $i = 1, \ldots, n$

and we call an $L(2,1)$-labeling c admissible if $c(x) \in T(x)$ for every $x \in V_G$. In any admissible $L(2,1)$-labeling, any pair of vertices must get distinct labels since G has diameter two. Moreover, as the vertices v_i^t and v_i^u are adjacent, they must be assigned labels that are at least two apart.

Hence $c(v_i^s) = s_i$, $c(v_i^t) = t_i$, and $c(v_i^u) = u_i$ is a one-to-one correspondence between admissible $L(2,1)$-labelings of G and systems of special representatives for $S_1, \ldots, S_m$ (the choice of $c(v_0) = 0$ or λ does not interfere with the labels of the remaining vertices). The graph G has clearly treewidth two. We will further design a collection of gadgets that will force the desired lists on the vertices of the graph G.

3.2 Labels of neighbors of vertices of large degrees The following simple observation will be used repeatedly in our arguments. Let v be a vertex whose two neighbors w and w' have degree $\lambda - 1$, and let c be an $L(2,1)$-labeling of span λ. Denote $S = c(N(w) \setminus \{v\})$ the set of labels used on the neighbors of w other than v. Since w and w' have the maximum possible degree, they are assigned labels 0 and λ, and hence $c(v) \in [2, \lambda - 2] \setminus S$.

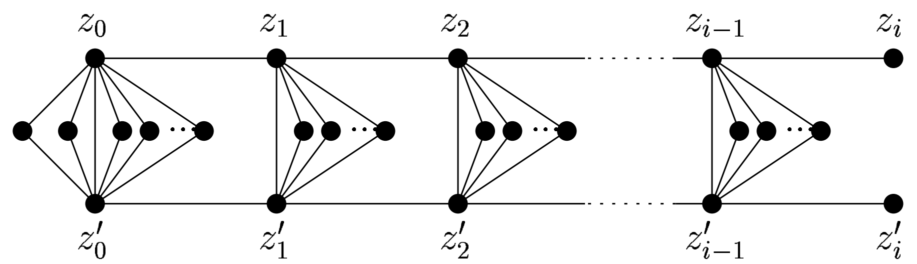

Figure 2: Construction of the graph H_i.

3.3 The crucial gadget For every $i \in [1, \frac{\lambda-1}{2}]$, we construct the graph H_i with nonadjacent vertices z_i, z'_i of degree one inductively as follows.

1) H_0 is the cycle of length four and z_0, z'_0 are two nonadjacent vertices (of degree two).

2) To construct H_{i+1}, we take the graph H_i and

- insert the edge (z_i, z'_i),
- insert two new vertices z_{i+1}, z'_{i+1} and edges $(z_i, z_{i+1}), (z'_i, z'_{i+1})$,
- insert $\lambda - 5$ new common neighbors of z_i and z'_i.

(See Fig. 2 for an example.) Then H_i is a series-parallel graph whose number of vertices is polynomial in i and n (precisely, $|V_{H_i}| = i(\lambda - 3) + 4$). It has the following crucial property.

Lemma 3.1. *For every $i \geq 1$, in any $L(2,1)$-labeling of H_i of span λ, the vertices $z_{i-1}, z_i, z'_{i-1}, z'_i$ are assigned (in this order) labels $i-1, \lambda-i, \lambda-i+1, i$ or $\lambda - i + 1, i, i-1, \lambda - i$.*

Proof. We prove the statement by induction on i. Let c be an $L(2,1)$-labeling of H_i of span λ.

1) For $i = 1$, observe that since z_0 and z'_0 have degree $\lambda - 1$, they must be assigned labels 0 and λ, or vice versa. Their $\lambda - 3$ common neighbors are assigned distinct labels forming the interval $[2, \lambda - 2]$ and hence $\{c(z_1), c(z'_1)\} = \{1, \lambda - 1\}$.

2) By induction hypothesis, $\{c(z_{i-1}), c(z'_{i-1})\} = \{i - 1, \lambda - i + 1\}$ and $\{c(z_i), c(z'_i)\} = \{i, \lambda - i\}$. These two vertices have further $\lambda - 5$ common neighbors that could be assigned only the labels forming the set $[0, i - 2] \cup [i + 2, \lambda - i - 2] \cup [\lambda - i + 2, \lambda]$. It is therefore easy to conclude that the two triples $(c(z_{i-1}), c(z'_i), c(z_{i+1}))$ and $(c(z'_{i-1}), c(z_i), c(z'_{i+1}))$ could be only the two consecutive triples $(i - 1, i, i + 1)$ and $(\lambda - i + 1, \lambda - i, \lambda - i - 1)$. $\square$

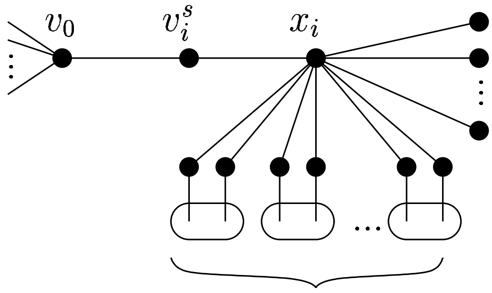

Figure 3: Forcing list S_i on the vertex v_i^s.

3.4 Forcing $T(v_0)$ Add $\lambda - 1 - 2n - m = 2n + 4 - m$ new neighbors to the vertex v_0. (We may assume $2n + 4 - m \ge 0$ since the SRL problem trivially has no system of distinct representatives if $2n + m > 2\lambda - 1$.) Then v_0 has degree $\lambda - 1$ and it can be assigned only labels 0 or λ by any $L(2,1)$-labeling of span λ.

3.5 Forcing $T(v_i^s)$ For each vertex $i \in [1, m]$, insert a new vertex x_i and make it adjacent to v_i^s. Further for each pair of labels l and $\lambda - l$ in the set $[2, \lambda - 2] \setminus S_i$, insert a new copy of the graph H_l and make x_i adjacent to the vertices z_l and z'_l of this new copy. Finally, add further new neighbors to the vertex x_i so that it has degree $\lambda - 1$ (see Fig. 3). It follows from the observation in 3.2 and Lemma 3.1 that the vertex v_i^s is now allowed to be assigned only a label from the set S_i as required.

3.6 Forcing $T(v_i^t)$ **and** $T(v_i^u)$ For each $i \in [1, n]$, insert vertices y_i, y'_i adjacent to v_i^t and v_i^u, respectively. Further take a copy of the graph H_{2i+1}, remove one common neighbor of z_{2i} and z'_{2i} and make y_i adjacent to z_{2i}, z_{2i+1} and y'_i to z'_{2i}, z'_{2i+1} of this copy. For each label $l \in [2, \frac{\lambda-1}{2}] \setminus \{2i, 2i+1\}$, insert two new copies of the graph H_l (the second copy is denoted by H_l^*) and connect both vertices z_l, z_l^* to y_i and both $z'_l, z_l^{*\prime}$ to y'_i. Finally, add three new neighbors to each vertex y_i, y'_i so that both have degree $\lambda - 1$ (see Fig. 4).

Suppose c is an $L(2,1)$-labeling of span λ. Since both y_i, y'_i have degree $\lambda - 1$ and are at distance 2 from v_0 of the same degree,they are both-assigned the same label, either 0 or λ. It also follows that in the copy H_{2i+1} the vertices z_{2i}, z'_{2i} behave as stated in Lemma 3.1, even if we removed one common neighbor (whose role was taken over by y_i and y'_i). Now according to observation in 3.2, the vertex v_i^t can be assigned only labels from $\{2i, 2i+1, \lambda - 2i - 1, \lambda - 2i\} \setminus \{c(z_{2i}), c(z_{2i+1})\}$ and similarly for

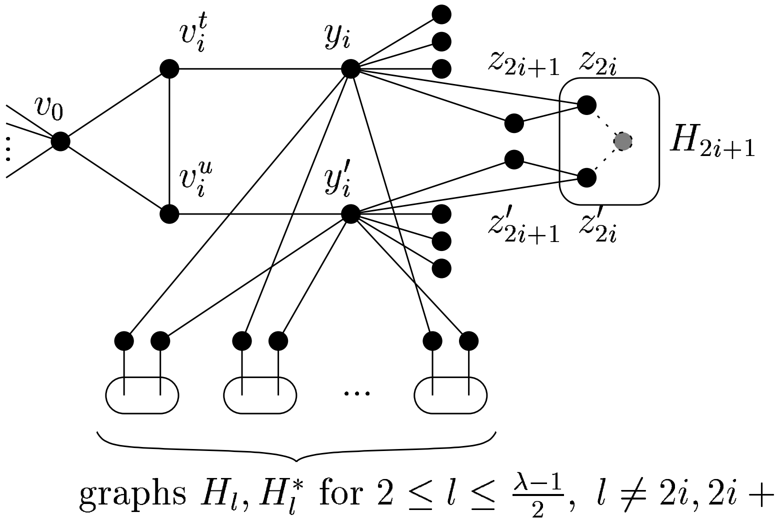

Figure 4: Forcing lists $\{2i, \lambda - 2i - 1\}$ and $\{2i+1, \lambda - 2i\}$ on the vertices v_i^t, v_i^u.

$c(v_i^u) \in \{2i, 2i+1, \lambda-2i-1, \lambda-2i\} \setminus \{c(z_{2i}'), c(z_{2i+1}')\}$. Since by Lemma 3.1 either $\{c(z_{2i}), c(z_{2i+1})\} = \{\lambda - 2i, 2i+1\}$ or $\{2i, \lambda-2i-1\}$, and respectively, $\{c(z_{2i}'), c(z_{2i+1}')\} = \{2i, \lambda - 2i - 1\}$ or $\{\lambda - 2i, 2i+1\}$, we get the desired admissible sets for both v_i^t and v_i^u (note here that the entire construction is symmetric with respect to vertices v_i^t and v_i^u).

By the above discussion, any $L(2,1)$-labeling of the resulting graph G' forces every vertex x of its subgraph G to be assigned labels from the list $T(x)$. During the construction of G' the distances between the original vertices of G were not changed, and hence any $L(2,1)$-labeling of G' restricted to G is an admissible $L(2,1)$-labeling for the lists $T(x), x \in V_G$.

The proof of the opposite implication (i.e., that any admissible $L(2,1)$-labeling of G can be extended to an $L(2,1)$-labeling of G') follows from the construction of all the gadgets and is straightforward.

Finally, observe that the size of G' is polynomial in the size of G (more precisely $|G'| = O(|G|^4)$, and also that all gadgets were constructed so that G' maintains treewidth two. This concludes the proof of Theorem 1.

4 $C(p,q)$-labelings of trees

Recall the notion of the cyclic metric, which depends on the span of the channel space. For $x, y \in [0, \lambda - 1]$, their distance is defined as $\rho_c(x,y) = \min\{|x-y|, \lambda - |x-y|\}$. A labeling $f : V_G \longrightarrow [0, \lambda-1]$ is called a $C(p,q)$-

labeling if $\rho_c(f(v), f(v')) \geq p$ for any edge $(v, v') \in E_G$, and $\rho_c(f(v), f(v')) \geq q$ for any two vertices $v, v' \in V_G$ of distance at most two in G. The minimum cyclic span λ for which a graph G admits a $C(p, q)$-labeling is denoted by $C_{(p,q)}(G)$.

Proposition 2 provides a slightly surprising fact that in the case of cyclic metric, the $C_{(p,q)}$ number of a tree is given by a closed formula and hence computable in linear time.

Also for the cyclic metric, we prove a dichotomy of the complexity of the C(2,1)-LABELING problem with respect to the treewidth of the input graph. However, it is worth noting that though the result is analogous to the case of $L(2, 1)$-labeling and the idea of the proof is similar (to reduce from the problem of special distant representatives via list labelings), the gadgets are constructed in a completely different way (they cannot be based on vertices of degree $\lambda - 1$ as in the case of $L(2, 1)$).

Before proving the theorem we first design a suitable gadget:

Lemma 4.1. *For any even λ and even integer $i \in [4, \frac{\lambda}{2}]$ there exists a series-parallel graph F_i of size $O(\lambda i)$ with two vertices z_i, z_i' of degree one, such that in any $C(2, 1)$-labeling f of F_i of span λ holds $\rho_c(f(z_i), f(z_i')) = i$.*

Proof. These gadgets are defined inductively: let F_2 be the path of length two with endvertices z_2 and z_2'.

The graph F_{i+2} is constructed from F_i by

- inserting the edge (z_i, z_i'),
- inserting two new vertices z_{i+2}, z_{i+2}' and edges $(z_2, z_{i+2}), (z_i', z_{i+2}')$,
- inserting $\lambda - 6$ new common neighbors of z_i and z_i'. (Consult Fig. 2 with the difference in H_1 and the numbers of inserted neighbors.)

As the initial vertices z_2 and z_2' are of degree $\lambda - 5$ they must be assigned labels of cyclic difference 2 and the labels of their $\lambda - 5$ common neighbors are uniquely determined as well as the labels of vertices z_4 and z_4'. If w.l.o.g $f(z_2) = 1$ and $f(z_2') = 3$ then $f(z_4) = 4$ and $f(z_4') = 0$.

It is straightforward to show by the same arguments as in Lemma 3.1 that for any i labels of z_{i-2}, z_i', z_{i+2} and z_{i-2}', z_i, z_{i+2}' are consecutive triples modulo λ and get the desired result. □

For simplicity we call vertices z_i and z_i' *the S (south) and N (north) poles* of F_i.

Observe that if we iterate the extension described in the lemma $\frac{\lambda}{2} - 1$ times, we get a graph whose poles allow labels of difference 2, but its size is $O(\lambda^2)$. In the future discussion we will recall this graph as F_2 instead of that used in the proof of Lemma 4.1.

Proof of Theorem 3. The proof is analogous to the proof of Theorem 1 and we show SRC $\propto$ $C(2,1)$-LABELING.

Let n and $S_1, \dots, S_m$ be the instance of the SRC problem. Observe that the transformation $x \to x' = 4x$ provides a bijection between the original SRC problem and its version where the aim is to find a collection of distinct integers $s'_1, s'_2, \dots, s'_m, t'_1, t'_2, \dots, t'_n, u'_1, u'_2, \dots, u'_n$ such that

- for all $j = 1, \dots, m : s'_j \in S'_j$, where all sets contain only multiples of 4,
- for all $i = 1, \dots, n : t'_i \in \{16n - 4i, 16n + 4i\}$,
- for all $i = 1, \dots, n : u'_i \in \{48n - 4i, 48n + 4i\}$,
- $\rho_c(t'_i, u'_i) = 32n - 8i$ for every $i = 1, 2, \dots, n$ and $\lambda = 64n$?

We construct the graph G according to the following plan (consult Fig. 5):

I. We first insert vertices $v_0, v^s_1, \dots, v^s_m, v^t_1, \dots, v^t_n, v^u_1, \dots, v^u_n$ where v_0 is adjacent to all others.

II. Build a chain of $2m$ copies of the graph F_{32n} such that the N pole of the i-th copy is merged with the S pole of the forthcoming copy and the resulting vertex is denoted by x_i. Identify the S pole of the first copy with v_0 and also join x_1 to v_0 by an edge.

III. For each $j = 1, \dots, m$ take every possible even $k \in \{2, 4, \dots, 32n-2\} \setminus S'_j$ and insert two copies of the graph F_k such that both S poles are identified with x_{2j} and both N poles are made adjacent to the vertex v^s_j.

IV. For each $i = 1, \dots, n$ join vertices v^t_i and v^u_i via gadget F_{32n-8i} such that v^t_i is merged with the S pole and v^u_i with the N pole.

V. Insert a copy of the graph F_{16n} and for each $i = 1, \dots, n$ insert two copies of the graph $F_{16n-12i}$ and two copies of $F_{16n+12i}$. Merge all S poles with the vertex v_0 and rename the N poles by y_{16n}, $y_{16n\pm12i}$ and $y'_{16n\pm12i}$.

VI. For each $i = 1, \dots, n$ join vertices v^t_i and y_{16n} via gadget F_{4i} such that v^t_i is merged with its S pole and y_{16n} with the other pole.

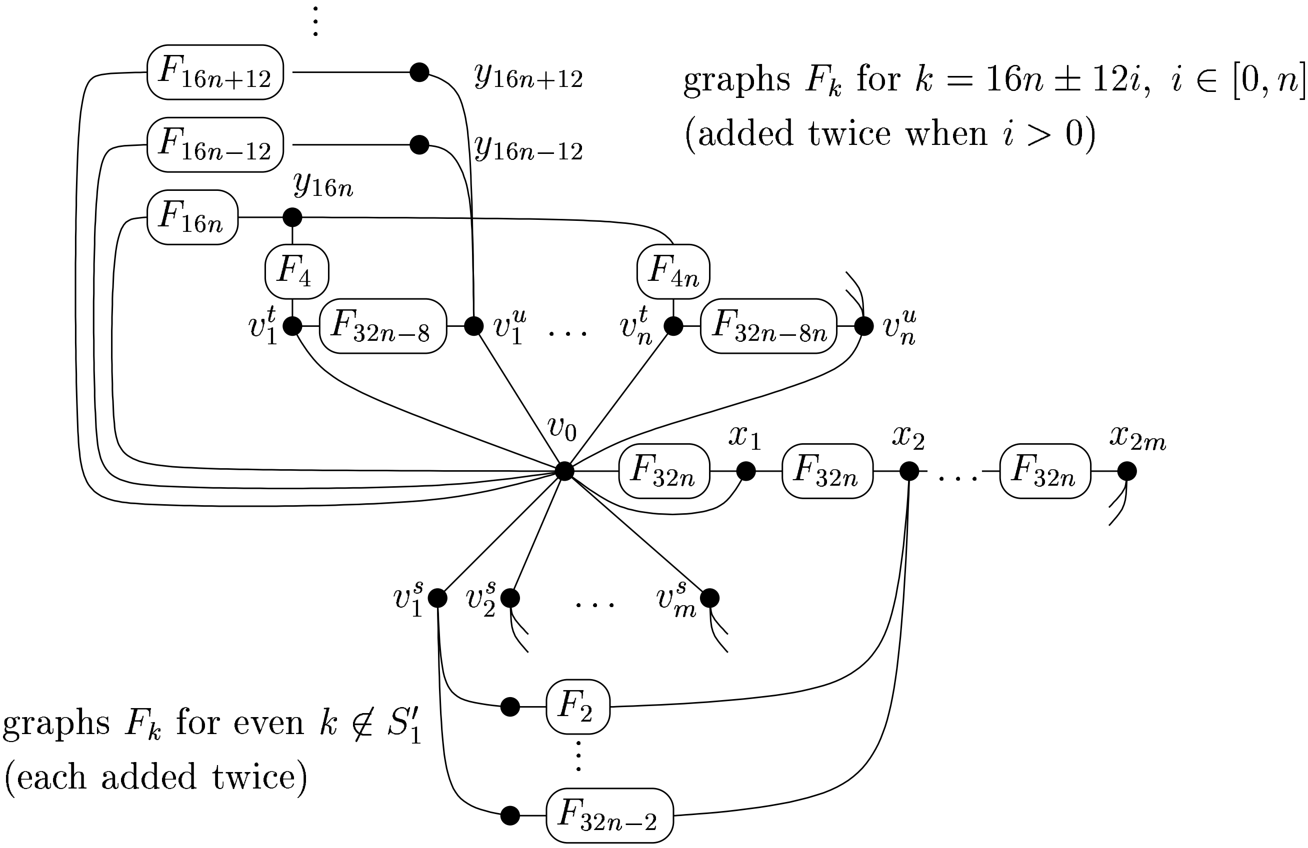

Figure 5: The graph G.

VII. Finally for each $i = 1, \ldots, n$ join v_i^u with the four vertices $y_{16n-12i}$, $y'_{16n-12i}$, $y_{16n+12i}$, $y'_{16n+12i}$.

Assume first that G allows a $C(2,1)$-labeling f of span $\lambda = 64n$. Without loss of generality we may assume $f(v_0) = 0$ and hence $f(x_{2j}) = 0$ for all $j = 1, \ldots, m$ (II.).

Each v_j^s have neighbors labeled $k, \lambda - k$ for all even $k \in \{2, 4, \ldots, 32n - 2\} \setminus S'_j$ due to (III.). Its neighbor v_0 is labeled by 0, and it cannot use label $32n$ since the vertex x_1 is forced label $32n$ and is at distance two from v_j^s. Hence only labels of S'_j remain feasible for v_j^s.

Due to symmetry of a $C(2,1)$-labeling we can further assume w.l.o.g. that $f(y_{16n}) = 16n$ and from (VI.) follows that only $16n \pm 4i$ are feasible labels for each v_i^t.

As $f(y_k) = k, f(y'_k) = \lambda - k$ or vice-versa, each v_i^u has forbidden labels $16n \pm 12i$ and $48n \pm 12i$ from (VII). Also due to (IV.,VI.) holds $f(v_i^u) \in f(v_i^t) \pm (32n - 8i) = \{48n \pm 4i, 48n \pm 12i\}$ and altogether only $48n \pm 4i$ remain feasible labels for v_i^u.

Finally due to (I.) all labels of $v_1^s, \ldots, v_n^u$ are distinct, hence $s_1, \ldots, u_n = \frac{1}{4}f(v_1^s), \ldots, \frac{1}{4}f(v_n^u)$ is a valid solution of the SRC problem.

In the opposite direction from any feasible system of representatives $s_1, \ldots, u_n$ we can derive a $C(2,1)$-labeling of G of span $\lambda = 64n$. Let the labels of vertices v_0, x_j, y_i, y_i' are settled as described above and labels of $v_1^s, \ldots, v_n^u$ are $s_1, \ldots, u_n$, each multiplied by 4.

To argue that this labeling can be extended to the entire graph G we note that as all copies of the graph F_k incident with any fixed vertex (especially v_0 and all x_{2j}) have distinct values of k, they invoke distinct *odd* labels on the adjoint vertices inside the gadgets F_k and cause no conflict with the other labels.

Finally observe that the resulting graph G is of treewidth 2. □

5 $(2,1)$-labelings of graphs of bounded treewidth

Given graphs G and H, an $H(2,1)$-labeling of G is a mapping $f : V_G \longrightarrow V_H$ such that adjacent vertices of G are mapped onto distinct nonadjacent vertices of H (i.e., distance of the target vertices is at least 2, measured in the target graph H) and vertices with a common neighbor in G are mapped onto distinct vertices of H (i.e., the distance of the target vertices is at least 1) [9]. This definition generalizes both the $L(2,1)$-labelings (when H is a path whose length equals the span of the labeling) and the $C(2,1)$-labelings (when H is a cycle whose length again equals the span). The computational complexity of this problem for fixed parameter graphs H was studied and many particular results were proven in [8]. The case when the span is also part of the input corresponds to the following decision problem:

$(2,1)$-LABELING

Input: Graphs G and H.

Question: Does G allow an $H(2,1)$-labeling?

Of course this problem is NP-complete for graphs G of treewidth two, since both L(2,1)-LABELING and C(2,1)-LABELING are its special cases. In this section we give a subtler separation of bounded width classes, namely in terms of pathwidth. Graphs of pathwidth one are caterpillars (trees obtained by pending any number of leaves to vertices of a path), and so the claim that $(2,1)$-LABELING is solvable in polynomial time for graphs G of pathwidth one (and arbitrary H) follows from our Theorem 4.

Proof of Theorem 4. The following algorithm is a straightforward extension of the algorithm for $L(2,1)$-labeling of trees of [4].

Given a tree T with m vertices, choose a leaf $r \in V_T$ and regard it as a root of T. For every edge $(u,v) \in E_T$ such that u is a child of v, denote by $T_{u,v}$ the subtree of T rooted in v and containing u and all its descendants. For every such edge and for every pair of vertices $x, y \in V_H$, we introduce a Boolean variable $\phi(u,v,x,y)$ which is true if and only if $T_{u,v}$ allows an $H(2,1)$-labeling f such that $f(u) = x$ and $f(v) = y$. Then T allows an $H(2,1)$-labeling if and only if $\phi(u,r,x,y) = \mathsf{true}$ for some vertices $x, y \in V_H$ (and u being the only child of the root r). The function ϕ can be computed by the following dynamic programming algorithm:

1. Set the initial values $\phi(u,v,x,y) = \mathsf{false}$ for all edges $(u,v) \in E_T$ and vertices $x, y \in V_H$.

2. If u is a leaf of T adjacent to its parent v, then set $\phi(u,v,x,y) = \mathsf{true}$ for all distinct nonadjacent vertices $x, y \in V_H$.

3. Suppose that ϕ is already calculated for all edges of $T_{u,v}$ except (u,v). Denote by $v_1, v_2, \dots, v_k$ the children of u. For all pairs of distinct nonadjacent vertices $x, y \in V_H$, construct the set system $\{M_1, M_2, \dots, M_k\}$, where
$$M_i = \{z \colon z \in V_H, z \neq y \text{ and } \phi(v_i, u, z, x) = \mathsf{true}\}$$
and set $\phi(u,v,x,y) = \mathsf{true}$ if the set system $\{M_1, M_2, \dots, M_k\}$ has a system of distinct representatives.

For the time analysis note that the recursive step requires, for each pair $x, y \in V_H$, time $O(nk)$ to construct the set system and time $O(k \cdot nk)$ for deciding if it has an SDR (e.g., by using the augmenting paths algorithm for a bipartite graph with at most nk edges and with k vertices in one bipartition class). Altogether the recursive step requires time $O(n^3k^2)$. If we denote by k_u the number of children of a nonleaf vertex $u \in V_T$, we have $\sum_{u \in V_T} k_u = m - 1$ (the number of edges of T), and hence the total running time is majorized by $O(\sum_{u \in V_T} n^3 k_u^2) = O(n^3 \sum_{u \in V_T} k_u^2) = O(n^3 (\sum_{u \in V_T} k_u)^2) = O(n^3 m^2)$. □

A *2-path* is a graph constructed from a triangle (say Δ_0) by consecutive augmentation of triangles so that each Δ_i shares one edge with the previously augmented Δ_{i-1}, while the third vertex of Δ_i is a vertex newly added in this step. A graph has pathwidth at most two if and only if it is a subgraph of a 2-path. In particular, a fan of triangles obtained from a path by adding a vertex adjacent to all vertices of the path, has pathwidth two.

Proof of Theorem 5. We reduce from HAMILTONIAN PATH which is well known to be NP-complete [13]. Given a graph H' with n vertices, let H be the disjoint union of the complement of H' and an isolated vertex x. Let G be obtained from a path of length $n-1$ on vertices $v_1, v_2, \ldots, v_n$ by adding a vertex w, which is adjacent to all v_i's. Then every $H(2,1)$-labeling f of G is an injective mapping from V_G to V_H (since G has diameter two), and f is necessarily bijective (since $|V_G| = |V_H|$). Without loss of generality $f(w) = x$, and hence $f(v_1), f(v_2), \ldots, f(v_n)$ is a Hamiltonian path in H', since $(v_i, v_{i+1}) \in E_G$ implies that $(f(v_i), f(v_{i+1})) \notin E_H$. The opposite implication is straightforward. □

6 Concluding remarks

We have fully characterized the computational complexity of (2,1)-distance constrained graph labelings in the case of linear and cyclic metrics in the channel space, with respect to the treewidth of the input graphs. Our results prove polynomial/NP-completeness dichotomy separating treewidth one from treewidth two, which is a rare phenomenon and has so far been known only for very few problems (namely the CUTWIDTH or MINIMUM LINEAR ARRANGEMENT which is polynomial for trees [24] while NP-hardness for graphs of treewidth two follows from [19]). With distance constrained labelings we have added a natural and important problem to this short list.

Let us remark that our main result is independent on the NP-completeness of the CHANNEL ASSIGNMENT problem, though both problems are related by the motivation in frequency assignment. The CHANNEL ASSIGNMENT is known NP-complete for graphs of treewidth three, but its complexity for treewidth two graphs is still open. The core of the NP-hardness of the two problems lies in different aspects of the problems and one does not straightforwardly follow from the other. On one hand, L(2,1)-LABELING relays to CHANNEL ASSIGNMENT by considering the second (distance) power of the input graph and assigning weights 2 to the original edges and 1 to the new ones. However, the graph constructed in this way will not have bounded treewidth. On the other hand, L(2,1)-LABELING involves only weights 2 and 1, while the NP-hardness of the CHANNEL ASSIGNMENT problem is based on large weights, the problem is not strongly NP-complete (it can be solved by dynamic programming algorithm in polynomial time if the weights are considered in unary encoding).

In the general (2,1)-LABELING problem, when both graphs come as parts of the input, we prove tight dichotomy with respect to pathwidth of the input (transmitters) graph. For both special metrics, $L(2,1)$ and $C(2,1)$, the complexity for graphs of bounded pathwidth is open.

To keep the paper well focused, we have stated most of the results for the simplest case of distance constraints (2,1). However, most of them can be extended to (p,q)- or at least $(p,1)$-labelings, see e.g. Proposition 2. It is known that L$(p,1)$-LABELING is polynomial for trees for every p (even the list and prelabeled versions), but when q does not divide p, the complexity of L(p,q)-LABELING for trees is open for all $q > 1$ (the list and prelabeled versions are known to be NP-complete [10]). To the contrary, C(p,q)-LABELING is polynomial for trees for all p, q as proven in our Theorem 2. Extension of our Theorem 4 to general $(p,1)$-LABELING of trees is trivial, since that follows by replacing H by its p-th distance power. An analog of Theorem 5 for general (p,q)-labelings can be proved by a more technical reduction.

References

[1] ARNBORG, S., AND PROSKUROWSKI, A. Characterization and recognition of partial k-trees *Congr. Numerantium 47* (1985) 69-75.

[2] BODLAENDER H. Polynomial algorithms for Graph Isomorphism and Chromatic Index on partial k-trees *J. Algorithms 11* (1990) 631-643.

[3] BODLAENDER H. A tourist guide through treewidth *Acta Cybern. 11* (1993) 1-21.

[4] CHANG, G. J., AND KUO, D. The $L(2,1)$-labeling problem on graphs. *SIAM Journal of Discrete Mathematics 9*, 2 (May 1996), 309–316.

[5] COURCELLE, B. The monadic second-order logic of graphs. I: Recognizable sets of finite graphs. *Inf. Comput. 85*, 1 (1990), 12–75.

[6] Fellows, M.R.; Kratochvíl, J.; Middendorf, M.; Pfeiffer, F. The complexity of induced minors and related problems *Algorithmica 13* (1995) 266-282.

[7] Fiala, J., Kloks, T., and Kratochvíl, J. Fixed-parameter complexity of λ-labelings *Discrete Applied Math. 113* (2001), 59-72.

[8] Fiala, J., and Kratochvíl, J. Complexity of partial covers of graphs In: Algorithms and Computation (P.Eades and T.Takaoka, eds.), Proceedings ISAAC 2001, Christchurch, December 2001, Lecture Notes in Computer Science 2223, Springer Verlag, Berlin Heidelberg 2001, pp. 537-549

[9] Fiala, J., and Kratochvíl, J. Partial covers of graphs *Discussiones Mathematicae Graph Theory 22* (2002) 89-99.

[10] Fiala, J., Kratochvíl, J., and Proskurowski, A. Distance constrained labelings of precolored trees. In *Theoretical Computer Science, 7th ICTCS '01, Torino* (2001), no. 2202 in Lecture Notes in Computer Science, Springer Verlag, pp. 285–292.

[11] Fiala, J., Kratochvíl, J., and Proskurowski, A. Systems of distant representatives. *Discrete Applied Mathematics 145*, 2 (2005), 306–316.

[12] Garey, M., Graham, R., Johnson, D., and Knuth, D. Complexity results for bandwidth minimization. *SIAM J. Appl. Math. 34* (1978), 477–495.

[13] Garey, M. R., and Johnson, D. S. *Computers and Intractability.* W. H. Freeman and Co., New York, 1979.

[14] Golovach, P. A. Systems of pairs of q-distant representatives and graph colorings. *Zap. nau. sem. POMI 293* (2002), 5–25. in Russian.

[15] Griggs, J. R., and Yeh, R. K. Labelling graphs with a condition at distance 2. *SIAM Journal of Discrete Mathematics 5*, 4 (Nov 1992), 586–595.

[16] Král, D., and Skrekovski, R. A theorem about the channel assignment problem *SIAM J. Discrete Math. 16*, No.3 (2003) 426-437.

[17] Liu, D., and Zhu, X. Circular Distance Two Labelings and Circular Chromatic Numbers. *Ars Combin. 69*, 4 (2003), 177–183.

[18] McDiarmid, C., and Reed, B. Channel assignment on graphs of bounded treewidth. *Discrete Math. 273*, 1-3 (2003), 183–192.

[19] Monien, B., and Sudborough, I. H. Min Cut is NP-Complete for Edge Weigthed Trees *Theor. Comput. Sci. 58*, No.1-3, (1988) 209-229.

[20] Roberts, F.S. private communication to J. Griggs.

[21] Robertson, N., and Seymour, P. Graph minors. I. Excluding a forest. *J. Comb. Theory, Ser. B 35* (1983), 39–61.

[22] Proskurowski, A. Separating subgraphs in k-trees: cables and caterpillars. *Discrete Math. 49* (1984), 275–285.

[23] van den Heuvel, J., Leese, R. A., and Shepherd, M. A. Graph labeling and radio channel assignment. *Journal of Graph Theory 29*, 4 (1998), 263–283.

[24] Yannakakis, M. A polynomial algoritm for the min-cut linear arrangements of trees *J. ACM 32* (1985) 950-988.