Computational Complexity of Generalized Domination: A Complete Dichotomy for Chordal Graphs

Petr Golovach[1,⋆] and Jan Kratochvíl[2,⋆⋆]

[1] Department of Informatics, University of Bergen, 5020 Bergen, Norway
petrg@ii.uib.no
[2] Department of Applied Mathematics and Institute for Theoretical Computer Science, Charles University, Prague, Czech Republic
honza@kam.mff.cuni.cz

Abstract. The so called (σ, ρ)-domination, introduced by J.A. Telle, is a concept which provides a unifying generalization for many variants of domination in graphs. (A set S of vertices of a graph G is called (σ, ρ)-*dominating* if for every vertex $v \in S$, $|S \cap N(v)| \in \sigma$, and for every $v \notin S$, $|S \cap N(v)| \in \rho$, where σ and ρ are sets of nonnegative integers and $N(v)$ denotes the open neighborhood of the vertex v in G.) It was known that for any two nonempty finite sets σ and ρ (such that $0 \notin \rho$), the decision problem whether an input graph contains a (σ, ρ)-dominating set is NP-complete, but that when restricted to chordal graphs, some polynomial time solvable instances occur. We show that for chordal graphs, the problem performs a complete dichotomy: it is polynomial time solvable if σ, ρ are such that every chordal graph contains at most one (σ, ρ)-dominating set, and NP-complete otherwise. The proof involves certain flavor of existentionality - we are not able to characterize such pairs (σ, ρ) by a structural description, but at least we can provide a recursive algorithm for their recognition. If ρ contains the 0 element, every graph contains a (σ, ρ)-dominating set (the empty one), and so the nontrivial question here is to ask for a maximum such set. We show that MAX-(σ, ρ)-domination problem is NP-complete for chordal graphs whenever ρ contains, besides 0, at least one more integer.

Keywords: Computational complexity, graph algorithms.

1 Introduction and Overview of Results

We consider finite undirected graphs without loops or multiple edges. The vertex set of a graph G is denoted by $V(G)$ and its edge set by $E(G)$. The open neighborhood of a vertex is denoted by $N(u) = \{v\colon uv \in E(G)\}$. A graph is chordal if it does not contain an induced cycle of length greater than three.

⋆ On leave from Department of Applied Mathematics, Syktyvkar State University, Syktyvkar, Russia. Most of the results were obtained during the research stay of the first author at DIMATIA Prague in 2006.

⋆⋆ Supported by the Czech Ministry of Education as Research Project No. 1M0545.

A. Brandstädt, D. Kratsch, and H. Müller (Eds.): WG 2007, LNCS 4769, pp. 1–11, 2007.

1.1 (σ, ρ)-Domination

Let σ, ρ be a pair of nonempty sets of nonnegative integers. A set of vertices of G is called (σ, ρ)-*dominating* if for every vertex $v \in S$, $|S \cap N(v)| \in \sigma$, and for every $v \notin S$, $|S \cap N(v)| \in \rho$. The concept of (σ, ρ)-domination was introduced by J.A. Telle [14,15] (and further elaborated on in [12,9]) as a unifying generalization of many previously studied variants of the notion of dominating sets (see [8] for an extensive bibliography on domination in graphs). In particular, $(\mathbb{N}_0,\mathbb{N})$-dominating sets are ordinary dominating sets, $(\{0\},\mathbb{N}_0)$-dominating sets are independent sets, $(\mathbb{N}_0,\{1\})$-dominating sets are efficient dominating sets, $(\{0\},\{1\})$-dominating sets are 1-perfect codes (or independent efficient dominating sets), $(\{0\},\{0,1\})$-dominating sets are strong stable sets, $(\{0\},\mathbb{N})$-dominating sets are independent dominating set, $(\{1\},\{1\})$-dominating sets are total perfect dominating set, or $(\{r\},\mathbb{N}_0)$-dominating sets are induced r-regular subgraphs ($\mathbb{N}$ and $\mathbb{N}_0$ denote the sets of positive and nonnegative integers, respectively).

We are interested in the complexity of the problem of existence of a (σ, ρ)-dominating set in an input graph, and we denote this problem by $\exists(\sigma, \rho)$-DOMINATION. It can be easily seen that if $0 \in \rho$, then the $\exists(\sigma, \rho)$-DOMINATION problem has a trivial solution $S = \emptyset$. So throughout the main part of the paper (and unless not explicitly stated otherwise) we suppose that $0 \notin \rho$.

1.2 Our Results

In view of the above given examples, it is not surprising that for any nontrivial combination of finite sets σ and ρ (considered as fixed parameters of the problem), $\exists(\sigma, \rho)$-DOMINATION is NP-complete [14]. It is then natural to pay attention to restricted graph classes for inputs of the problem. It was observed in [11] that for any pair of finite sets σ and ρ, the problem is solvable in polynomial time for interval graphs, but that it becomes NP-complete when restricted to chordal graphs (for some parameter sets σ and ρ). In particular, it was shown that for one-element sets $\sigma = \{p\}, \rho = \{q\}$, $\exists(\sigma, \rho)$-DOMINATION is polynomial time solvable if $q > 2p + 1$ and NP-complete if $q \leq p + 1$. We close this gap by showing that all the remaining cases are also polynomial time solvable. Moreover, we extend this polytime/NP-completeness dichotomy to any pair of finite sets σ, ρ by showing the following characterization:

Theorem A. *For finite sets σ, ρ, $\exists(\sigma, \rho)$-DOMINATION is polynomial time solvable for chordal graphs if every chordal graph has at most one (σ, ρ)-dominating set, and it is NP-complete otherwise.*

This theorem provides a full characterization and dichotomy, with both the polynomial time solvable and NP-complete cases including nontrivial and interesting samples (as we show by discussing some examples in Section 4). Dichotomy results are valued and intensively looked for (e.g., the classification of Boolean satisfiability by Schaefer [13], further dichotomy results for larger classes of the Constraint Satisfaction Problem by Bulatov et al. [2] paving the way to the utmost CSP dichotomy conjecture of Feder and Vardi [4], or several results for graph homomorphisms [10,3,6,5].) The characterization is nonconstructive in the sense that we are not able to provide a structural description of ambivalent (or non-ambivalent) pairs σ, ρ (we call a pair σ, ρ *ambivalent* if there exists a chordal graph containing two different (σ, ρ)-dominating sets), and there is indication that such a description will not be simple. Indeed, for any pair of σ and ρ, there are infinitely many chordal graphs to be checked if any of them, by chance, contains two different (σ, ρ)-dominating sets. Perhaps somewhat surprisingly we show that this fact can be overcome at least from the computational point of view:

Theorem B. *It can be decided in finite time (i.e., by a recursive algorithm) whether for a given pair of finite sets σ, ρ, there exists a chordal graph containing two different (σ, ρ)-dominating sets.*

The NP-hardness part of Theorem A is proved in Section 2 by a reduction from a variant of the Exact Cover problem. Its polynomial part is proved in Section 3 by providing an explicit dynamic programming algorithm. Theorem B is proved by providing an explicit upper bound on the minimum size of an ambivalent graph in Section 4. In Section 5, we discuss the case when $0 \in \rho$. As we have already mentioned, the $\exists(\sigma, \rho)$-domination problem is then trivial (the empty set is always (σ, ρ)-dominating), and the natural question here is the optimization variant. However, we show this is always a hard problem:

Theorem C. *Given a chordal graph graph G and a number k, it is NP-complete to decide if G contains a (σ, ρ)-dominating set of size at least k, provided σ, ρ are finite sets of nonnegative integers and $\rho \neq \{0\}$.*

Throughout the paper $n = |V(G)|$, $p_{\min} = \min \sigma, p_{\max} = \max \sigma, q_{\min} = \min \rho$ and $q_{\max} = \max \rho$, where G is the graph and σ, ρ the sets under consideration. In case of single-element sets σ or ρ, we write simply $p = p_{\min} = p_{\max}$ and $q = q_{\min} = q_{\max}$.

2 NP-Complete Cases

This section is devoted to the proof of the following theorem.

Theorem 1. *Let σ, ρ be finite sets of nonnegative integers, $0 \notin \rho$. If there is a chordal graph with at least two different (σ, ρ)-dominating sets, then the $\exists(\sigma, \rho)$-domination problem is NP-complete for chordal graphs.*

2.1 An Auxiliary Complexity Lemma

We are going to reduce from a special variant of the Cover by triples problem (or Exact Cover)(see [7]).

Let r be a positive integer. An instance of the Cover by no more than r triples is a pair (X, M), where X is a nonempty finite set and M is a set of triples of elements of X. We ask about the existence of a set $M' \subset M$ such that every element of X belongs to at least one and to at most r triples of M'. Such a set we call a *cover of X by no more than r triples.* For space limitations the proof of the following auxiliary lemma is omitted.

Lemma 1. *For every fixed $r \geq 1$, the* Cover by no more than r triples *problem is NP-complete.*

2.2 The Forcing Gadget

Our next step of the proof is the construction of a gadget which "enforces" on a given vertex the property of "not belonging to any (σ, ρ)-dominating set".

It is known (cf. [11]) that if $q_{\min} \geq 2p_{\max}+2$, then every chordal graph contains at most one (σ, ρ)-dominating set. Hence we assume that $q_{\min} \leq 2p_{\max}+1$. We construct a rooted graph F as follows.

Suppose first that $q_{\min} \leq p_{\max}+1$. We start with a complete graph $K_{p_{\max}+1}$ with vertices $u_1, u_2, \ldots, u_{p_{\max}+1}$. Let $\{S_1, S_2, \ldots, S_t\}$ be a set of $q_{\min}$-tuples which covers the set $\{u_1, u_2, \ldots, u_{p_{\max}+1}\}$ (i.e., each u_j belongs to at least one S_i). For every $i = 1, 2, \ldots, t$, we add $q_{\max}+1$ new vertices $v_1^{(i)}, v_2^{(i)}, \ldots, v_{q_{\max}+1}^{(i)}$ and connect them to all vertices of S_i by edges.

If $q_{\min} > p_{\max}+1$, the construction is slightly different. We again start with a complete graph $K_{p_{\max}+1}$ with vertices $u_1, u_2, \ldots, u_{p_{\max}+1}$. We add $q_{\max}+1$ new vertices $v_1, v_2, \ldots, v_{q_{\max}+1}$ and $q_{\max}+1$ copies of $K_{p_{\max}+1}$, say $Q_1, Q_2, \ldots, Q_{q_{\max}+1}$, and connect every v_j by edges to all vertices $u_1, u_2, \ldots, u_{p_{\max}+1}$ and to $q_{\min}-p_{\max}+1$ vertices of the corresponding Q_j.

In both cases the vertex u_1 is the root of F.

Lemma 2. *The graph F has at least one (σ, ρ)-dominating set, and for every (σ, ρ)-dominating set S in F, $u_1, u_2, \ldots, u_{p_{\max}+1} \in S$. Moreover, if F is an induced subgraph of a graph F' such that u_1 is the only vertex of F adjacent to vertices of $F' \setminus F$, then the vertices of $F' \setminus F$ that are adjacent to u_1 do not belong to any (σ, ρ)-dominating set in F'.*

Proof. Suppose that $q_{\min} \leq p_{\max}+1$. Obviously $\{u_1, u_2, \ldots, u_{p_{\max}+1}\}$ is a (σ, ρ)-dominating set in F. For the second statement, assume that S is a (σ, ρ)-dominating set in F and $u_i \notin S$ for some i. Let S_j be a $q_{\min}$-tuple which contains u_i. It is readily seen that $v_1^{(j)}, v_2^{(j)}, \ldots, v_{q_{\max}+1}^{(j)} \in S$. But then u_i is adjacent to at least $q_{\max}+1$ vertices of S, a contradiction.

If $q_{\min} > p_{\max}+1$, the proof of the second statement is similar. For the first part, note that the vertices $u_1, u_2, \ldots, u_{p_{\max}+1}$ and all vertices of the added cliques Q_j form a (σ, ρ)-dominating set.

For the last statement, note that we have proved that in both cases u_1 is in S and has $p_{\max}$ neighbors in S, for any (σ, ρ)-dominating set S in F, but the argument survives for any (σ, ρ)-dominating set in F' as well. □

2.3 The Reduction

Let H be a graph which has at least two different (σ, ρ)-dominating sets $S, \widetilde{S}$. We choose a vertex $u \in S \div \widetilde{S}$, where $\div$ denotes the symmetric difference of sets, and pronounce u the root of H. Let $k = \max\{i \in \mathbb{N}_0 : i \notin \rho, i+1 \in \rho\}$. Since $0 \notin \rho$, k is correctly defined.

Let a set $X = \{x_1, x_2, \ldots, x_n\}$ and a set $M = \{t_1, t_2, \ldots, t_m\}$ of triples on X be given as an instance of Cover by no more than r triples for $r = q_{\max} - k > 0$.

We start the construction of a graph G with a complete graph K_n with vertices $x_1, x_2, \ldots, x_n$. For every triple $t_i = \{x_a, x_b, x_c\}$, a copy H_i of the graph H with root u_i is added, and u_i is connected by edges to x_a, x_b, x_c. If $k = 0$, we further add $q_{\max}$ copies of the graph F with roots $v_1, v_2, \ldots, v_{q_{\max}}$, add a new extra vertex y, and join y with $x_1, x_2, \ldots, x_n$ and $v_1, v_2, \ldots, v_{q_{\max}}$ by edges. If $k > 0$, then k copies of F with roots $v_1, v_2, \ldots, v_k$ are added, and vertices $v_1, v_2, \ldots, v_k$ are connected with $x_1, x_2, \ldots, x_n$ by edges.

We claim that the graph G constructed in this way has a (σ, ρ)-dominating set if and only if (X, M) allows a cover by no more than r triples. Since the graphs H and F depend only on σ and ρ, G has $O(n + m)$ vertices, our reduction is polynomial and the proof will be concluded.

Suppose first that G has a (σ, ρ)-dominating set S. Let $M' = \{t_i \in M : u_i \in S\}$. If $k = 0$, then $y \notin S$ and $v_1, v_2, \ldots, v_{q_{\max}} \in S$ by Lemma 2. Hence $x_1, x_2, \ldots, x_n \notin S$. Since $0 \notin \rho$, for every $i = 1, 2, \ldots, n$, the vertex x_i has at least one S-neighbor in the set $\{u_1, u_2, \ldots, u_m\}$, but no more than $r = q_{\max}$ such neighbors. So M' is a cover of X by no more than r triples.

If $k > 0$, then $v_1, v_2, \ldots, v_k \in S$ and $x_1, x_2, \ldots, x_n \notin S$ by Lemma 2 again. Since $k \notin \rho$, for every $i = 1, 2, \ldots, n$, the vertex x_i has at least one S-neighbor in the set $\{u_1, u_2, \ldots, u_m\}$, but no more than $r = q_{\max} - k$ such neighbors. Hence again, M' is a cover of X by no more than r triples.

Suppose now that $M' \subseteq M$ is a cover of X by no more than r triples. For every $i = 1, 2, \ldots, m$, we choose a (σ, ρ)-dominating set S_i in H_i such that $u_i \in S_i$ if and only if $t_i \in M'$. Let $S'_1, S'_2, \ldots$ be (σ, ρ)-dominating sets in the copies of F. Since $\{k+1, k+2, \ldots, q_{\max}\} \subseteq \rho$, $S = S_1 \cup S_2 \cup \cdots \cup S_n \cup S'_1 \cup S'_2 \cup \ldots$ is a (σ, ρ)-dominating set in G.

3 The Polynomial Cases

In this section we prove the complementary part of Theorem A by presenting a polynomial time algorithm that decides the existence of a (σ, ρ)-dominating set in a chordal graph, provided the parameters σ and ρ are such that every chordal graph contains at most one (σ, ρ)-dominating set. It is perhaps of some interest that our algorithm can be formulated in a general way so that it is based only on the promise of a unique solution. On the contrary, in many situations the assumption of uniqueness of the solution does not help.

In fact we present two algorithms in this section. In the first subsection we give the general algorithm, and in the latter one we deal with a special case of one-element set σ. The running time of the second algorithm is much better and moreover, this algorithm explicitly closes the gap between polynomial and NP-complete cases for single-element parameter sets left open in [11].

3.1 The General Algorithm

In this subsection it is assumed that σ and ρ are such that every chordal graph contains no more than one (σ, ρ)-dominating set. The algorithm uses dynamic programming and is based on the clique-decomposition of the input graph.

Let $\mathcal{K}$ be the set of all maximal cliques of an input chordal graph G, and let T be a clique tree of G, i.e., $V(T) = \mathcal{K}$ and for every $u \in V(G)$, the subgraph of T induced by $\{K \in \mathcal{K} : u \in K\}$ is connected. It is well known (see, for example, [1]) that a clique tree of a chordal graph graph is not unique, but can be constructed in linear time. We choose a clique $R_0 \in \mathcal{K}$ and consider the clique tree T rooted in R_0. This induces a parent-child relation in the tree, in which all vertices are descendants of the root. For any clique $R \in \mathcal{K}$, we denote by T_R the subtree of T rooted in R and containing all descendants of R, and we denote by G_R the subgraph of G induced by the vertices contained in the cliques of $V(T_R)$.

The key idea of the algorithm is the fact that every clique $R \in \mathcal{K}$ contains at most $p_{\max} + 1$ vertices of any (σ, ρ)-dominating set, and hence for every clique we can list all possible intersections with a solution set S in polynomial time. We need to keep track of how many S-neighbors these vertices have. Towards this end we build the following array. Let $R \in \mathcal{K}$ and let $X = \{x_1, x_2, \ldots, x_r\}$ be an ordered subset of R, $0 \le r \le p_{\max} + 1$ (X can also be empty). Further let $P = (p_1, p_2, \ldots, p_r)$ be a sequence of nonnegative integers, $p_i \le p_{\max}$ for $i = 1, 2, \ldots, r$. For each such triple R, X, P, our algorithm constructs a set $S(R, X, P) \subseteq V(G_R)$ which satisfies

- $S(R, X, P) \cap R = X$,
- $|N(x_i) \cap S(R, X, P)| = p_i$ for $i = 1, 2, \ldots, r$,
- for every $v \in V(G_R) \setminus R$, $|N(v) \cap S(R, X, P)| \in \begin{cases} \sigma & \text{if } v \in S(R, X, P), \\ \rho & \text{if } v \notin S(R, X, P); \end{cases}$

(i.e., $S(R, X, P)$ is a candidate for $S \cap V(G_R)$) or $S(R, X, P) =$NIL if we can deduce that no such set can be extended to a solution S for the entire G. The details of the algorithm will appear in the full version of the paper. The recursive step is technical but straightforward. The crucial fact is that for each triple R, X, P, we store at most one candidate set, which follows from the following lemma.

Lemma 3. *If S_1 and S_2 are distinct subsets of $V(G_R)$ satisfying the candidate conditions for the same triple R, X, P, then none of them can be extended to a (σ, ρ)-dominating set S in the entire graph G.*

Proof. Suppose S_1 can be extended to a (σ, ρ)-dominating set S. Then S and $(S \setminus S_1) \cup S_2$ are two distinct (σ, ρ)-dominating sets in $G[V(G) \setminus (R \setminus X)]$, which is a contradiction to the assumption that every chordal graph contains at most one (σ, ρ)-dominating set. □

The algorithm can be implemented to run in time $O(n^{p_{\max}^2 + 2p_{\max} + 3})$. Hence we have proved the polynomial part of Theorem A:

Theorem 2. *If σ and ρ are finite sets such that every chordal graph contains at most one (σ, ρ)-dominating set, then the* $\exists(\sigma, \rho)$-DOMINATION *problem is solvable in polynomial time for chordal graphs.*

3.2 Single-Element Sigma

In the case when the set σ contains only one element, we are able to design a more efficient greedy algorithm. This algorithm uses the simple structure of (σ, ρ)-dominating sets in such a case.

Lemma 4. *Let $\sigma = \{p\}$, let ρ be arbitrary and suppose that S is a (σ, ρ)-dominating set in a chordal graph G. Then S is the union of disjoint cliques of size $p + 1$, and vertices of different cliques are nonadjacent.*

Proof. Let $G[S]$ be the subgraph of G induced by S. This graph is chordal, so it has a simplicial vertex. The closed neighborhood of this vertex is a clique of size $p + 1$, and this clique is the vertex set of one component of $G[S]$. By repeating these arguments we prove that all components of $G[S]$ are induced by cliques of size $p + 1$. □

Though we use the following observation for the case of single-element σ, we state it in a more general form:

Lemma 5. *Suppose that $p_{\max} + 2 \leq q_{\min}$. Let S be a (σ, ρ)-dominating set in G. Then all simplicial vertices of G belong to S.*

Proof. Let v be a simplicial vertex of G. If $v \notin S$, then $|N(v) \cap S| \in \rho$, and since $p_{\max} + 2 \leq q_{\min}$, $|N(v) \cap S| \geq p_{\max} + 2$. So S contains a clique of size $p_{\max} + 2$, a contradiction. □

The core observation for our algorithm is the following lemma, which is a straightforward corollary of Lemmas 4 and 5.

Lemma 6. *Let $\sigma = \{p\}$ and let $p + 2 \leq q_{\min}$. Let S be a (σ, ρ)-dominating set in a chordal graph G. Further let T be a clique tree of G, let X be a leaf of T, and Y the neighbor of X in T. Then*

- $|X \setminus Y| \leq p + 1$,
- *if* $|X \setminus Y| = p + 1$, *then* $X \setminus Y \subseteq S$ *and* $(Y \cap X) \cap S = \emptyset$,
- *if* $|X \setminus Y| < p + 1$, *then* $(Y \setminus X) \cap S = \emptyset$.

Given a chordal graph, our algorithm first builds a clique tree and then consecutively reduces it by deleting vertices which must or must not belong to any (σ, ρ)-dominating set. In the final step it is necessary to check whether the only candidate (if any) for S is really a (σ, ρ)-dominating set. The technical details will again appear in the full version of the paper. We only note that with some extra care the algorithm can be designed so that in each reduction step, a clique tree of the reduced graph can be easily derived from the clique tree of the previous one. Thus we can claim:

Theorem 3. *If $\sigma = \{p\}$ and $p + 2 \leq q_{\min}$, then the* (σ, ρ)-DOMINATION *problem can be solved in time $O(n^2)$.*

4 Uniqueness of (σ, ρ)-Dominating Sets

It would be most desirable to have a full classification of the pairs of parameter sets σ, ρ for which there exist chordal graphs with two different (σ, ρ)-dominating sets. Such a classification is not currently known and perhaps not easy to obtain. In the first subsection of this section we summarize the known results in this direction. A positive result is proven in the second subsection. We show a bound on the size of a minimal chordal graph containing two different (σ, ρ)-dominating sets, thus showing that the existence of such a graph can be decided by a finite algorithm.

Recall that we call a pair (σ, ρ) *ambivalent* if there exists a graph containing at least two different (σ, ρ)-dominating sets. Such a graph will be called (σ, ρ)-*ambivalent*.

4.1 On the Way to Classification

First observation about the uniqueness of a (σ, ρ)-dominating set in a chordal graph was made in [11]. More cases are covered by the following theorem, but the picture is far from being complete. Fully characterized are the cases of $\sigma = \{p\}$ and $\sigma = \{0, p_{\max}\}$.

Theorem 4. *The following table presents examples of ambivalent and non-ambivalent pairs of σ and ρ:*

ambivalent	*non-ambivalent*
$q_{\min} \leq p_{\max} + 1$	$q_{\min} \geq 2p_{\max} + 2$
$\exists i : \{i, i+1\} \subseteq \sigma,\ q_{\min} \leq 2p_{\max} + 1$	$\sigma = \{p\},\ q_{\min} \geq p + 2$
$\sigma = \{0, p_{\max}\},\ q_{\min} \leq p_{\max} + 2$	$\sigma = \{0, p_{\max}\},\ q_{\min} \geq p_{\max} + 3$

Proof. Will appear in the full version. □

4.2 Deciding the Ambivalence

The main goal of this subsection is to prove Theorem B. We do so by proving an upper bound on the number of vertices of any minimum chordal graph containing two different (σ, ρ)-dominating sets, in terms of $p_{\max}$ and $q_{\max}$.

Theorem 5. *Let σ, ρ be finite sets of nonnegative integers, $0 \notin \rho$. Suppose that G is a minimum chordal (σ, ρ)-ambivalent graph. Then*

- *for every maximal clique K of G, $|K| \leq 2p_{\max} + 2$,*
- *for every vertex $v \in V(G)$, $\deg v \leq \max\{2p_{\max}, p_{\max} + q_{\max}\}$,*
- *the diameter of G is $O(p_{\max}^{2p_{\max}+2q_{\max}+7})$.*

Proof. Will appear in the full version. □

Since every graph of maximum degree Δ and diameter d has at most Δ^{d+1} vertices, we have proven the following corollary and hence also Theorem B.

Corollary 1. *Let σ, ρ be finite sets of nonnegative integers, $0 \notin \rho$. Then the size of every minimum (σ, ρ)-ambivalent chordal graph is bounded by a function of $p_{\max}$ and $q_{\max}$ and the existence of such a graph can be tested algorithmically by a finite procedure.*

5 MAX-(σ, ρ)-Domination

So far we have been considering the question of existence of (σ, ρ)-dominating sets. One could also pose the optimization questions, i.e., asking for the sizes of minimum or maximum (σ, ρ)-dominating sets. Since optimization problems are at least as difficult as the existence ones, and since the polynomial part of our Theorem A is based on uniqueness of the solution, our results translate directly to the optimization variants. Namely, if $0 \notin \rho$, then both MIN-(σ, ρ)-DOMINATION and MAX-(σ, ρ)-DOMINATION problems are NP-hard when restricted to chordal graphs for ambivalent (σ, ρ) and polynomial time solvable for the non-ambivalent pairs.

If $\rho = \{0\}$, the only possible (σ, ρ)-dominating sets in a connected graph G are $S = \emptyset$ and $S = V(G)$. The latter is the maximum (σ, ρ)-dominating set if $\deg v \in \sigma$ for every $v \in V(G)$, otherwise $S = \emptyset$ is the only (and hence also the maximum) (σ, ρ)-dominating set in G. This is, however, the only polynomially solvable case, as Theorem C claims. The rest of this section is devoted to its proof.

5.1 Proof of Theorem C

We begin with an auxiliary construction. Let F consist of $q_{\max}$ copies of the complete graph $K_{p_{\max}+1}$, say $Q_1, Q_2, \ldots, Q_{q_{\max}}$, and one extra vertex r, the root of F, which is adjacent to exactly one vertex from each Q_i. The following technical lemma is straightforward.

Lemma 7. *The set $S = V(Q_1) \cup V(Q_2) \cup \cdots \cup V(Q_{q_{\max}})$ is a maximum (σ, ρ)-dominating set in F, and it has cardinality $q_{\max}(p_{\max} + 1)$. Moreover, suppose that a graph F' is created by uniting F with some graph (with different vertices) and joining the root of F to some new vertices $u_1, u_2, \ldots, u_s$. If S' is a (σ, ρ)-dominating set in F', $r \notin S'$, and $u_i \in S'$ for some i, then $|V(F) \cap S'| < q_{\max}(p_{\max} + 1)$.*

Now we prove Theorem C by a reduction from the EXACT h-COVER problem, whose is input is a pair (X, M), where $X = \{x_1, x_2, \ldots, x_n\}$ is a finite set and $M = \{t_1, t_2, \ldots, t_m\}$ is a set of triples on X, and the question is if M contains a subsystem $M' \subset M$ such that every element of X belongs to exactly h triples of M'. This problem is NP-complete for every fixed $h > 0$ (cf. e.g., [11]). For our reduction, we use $h = q_{\max}$. For a given instance (X, M), we may assume without loss of generality that $nq_{\max} = 3l$ and $l \leq m$.

We start the construction of a graph G with a complete graph K_n with vertices $x_1, x_2, \ldots, x_n$. For every triple $t_i = \{x_a, x_b, x_c\}$, a copy H_i of the complete graph $K_{p_{\min}+1}$ is added, and one vertex of this graph is connected by edges to x_a, x_b, x_c. We further add $s = (m-l)(p_{\min}+1) + 2p_{\max} + 2$ copies $F_1, \ldots, F_s$ of the graph F, with their roots $r_1, r_2, \ldots, r_s$ being adjacent to all $x_1, x_2, \ldots, x_n$. This graph G has $m(p_{\min}+1)+n+s(q_{\max}(p_{\max}+1)+1)$ vertices and it is constructed from (X, M) in polynomial time. We claim that G contains a (σ, ρ)-dominating set of size $\geq k = l(p_{\min}+1) + sq_{\max}(p_{\max}+1)$ if and only if (X, M) contains an exact $q_{\max}$-cover.

Let first M' be a $q_{\max}$-cover of X. Clearly, $|M'| = l$, and it is straightforward to check that $S = \bigcup_{i:\, t_i \in M'} V(H_i) \cup \bigcup_{i=1}^{s} (V(F_i) \setminus \{r_i\})$ is a (σ, ρ)-dominating set in G of cardinality k.

Assume now that S is a (σ, ρ)-dominating set in G, and $|S| \geq k$. Suppose that some vertex x_j is in S. Then S can contain no more than $m(p_{\min}+1)$ vertices of the graphs H_i, and no more than $p_{\max}+1$ vertices from the set $\{x_1, x_2, \ldots, x_n\}$. Also at least $s - p_{\max}$ vertices from $\{r_1, r_2, \ldots, r_s\}$ do not belong to S. So, according to the preceding lemma, $|S| \leq m(p_{\min}+1)+p_{\max}+1+p_{\max}q_{\max}(p_{\max}+1)+(s-p_{\max})(q_{\max}(p_{\max}+1)-1) = m(p_{\min}+1)+2p_{\max}+1+sq_{\max}(p_{\max}+1)-(m-l)(p_{\min}+1) - 2p_{\max} - 2 = l(p_{\min}+1) + sq_{\max}(p_{\max}+1) - 1 < k$. So, none of the vertices $x_1, x_2, \ldots, x_n$ belongs to S. Note that in this case $V(H_i) \subset S$ or $V(H_i) \cap S = \emptyset$ for all $i = 1, 2, \ldots, m$, and vertices of no more than l graphs H_i belong to S. Since S can contain no more than $q_{\max}(p_{\max}+1)$ vertices from every graph F_i, and $|S| \geq k$, vertices of exactly l graphs H_i are included to S. Every vertex x_j can have no more than $q_{\max}$ adjacent vertices from S. Hence each x_j is adjacent to exactly $q_{\max}$ vertices from H_i's and the set $M' = \{t_i \colon V(H_i) \subset S\}$ is a $q_{\max}$-cover of X.

6 Concluding Remarks and Open Problems

The complete classification of ambivalent pairs (σ, ρ) remains the first and main open problem. We believe that it is an interesting combinatorial problem by itself, and that it deserves attention. Perhaps it is impossible to formulate simple necessary and sufficient conditions for the general problem, but it would be interesting to obtain a complete solution at least for some special cases. For example for two-element sets $\sigma = \{p_1, p_2\}$ (it seems that cardinality of σ is more important).

A related complexity question is if the ambivalence of (σ, ρ) can be tested in polynomial time.

Another interesting question is a fixed parameter tractability of the $\exists(\sigma, \rho)$-domination. If the maximal value of σ is supposed to be the parameter, then the Theorem 3 shows that this problem is in FPT for $|\sigma| = 1$ and $p + 2 \leq q_{\min}$ (in fact our algorithm is polynomial in p and n). On the other hand, our general algorithm from Subsection 3.1 has the parameter $p_{\max}$ in the exponent of the running time, and hence is not FPT-algorithm. Fixed parameter tractability (or intractability) of the general case remains an open problem. Also it would be interesting to consider the problem parametrized by the size of the (σ, ρ)-dominating set.

References

1. Blair, J.R.S., Peyton, B.W.: An introduction to chordal graphs and clique tree. In: George, J.A., Gilbert, J.R., Liu, J.W.H. (eds.) Sparse Matrix Computation: Graph Theory Issues and Algorithms. IMA Volumes in Mathematics and its Applications, vol. 56, pp. 1–30. Springer, Heidelberg (1993)
2. Bulatov, A.A., Jeavons, P., Krokhin, A.A.: Classifying the complexity of constraints using finite algebras. SIAM J. Comput. 34(3), 720–742 (2005)
3. Feder, T., Hell, P., Huang, J.: Bi-arc graphs and the complexity of list homomorphisms, J. Graph Theory 42, 61–80 (2003)
4. Feder, T., Vardi, M.Y.: The computational structure of momotone monadic SNP and constraint satisfaction: A sudy through datalog and group theory. SIAM Journal of Computing 1, 57–104 (1998)
5. Fiala, J., Kratochvil, J.: Locally injective graph homomorphism: Lists guarantee dichotomy. In: Fomin, F.V. (ed.) WG 2006. LNCS, vol. 4271, pp. 15–26. Springer, Heidelberg (2006)
6. Fiala, J., Paulusma, D.: A complete complexity classification of the role assignment problem. Theoretical Computer Science 1 349, 67–81 (2005)
7. Garey, M., Johnson, D.: Computers and intractability. W.H.Freeman, New York (1979)
8. Haynes, T., Hedetniemi, S., Slater, P.: Domination in Graphs: The Theory. Marcel Dekker, New York (1997)
9. Heggernes, P., Telle, J.A.: Partitioning graphs into generalized dominating sets. Nordic J. Comput. 5, 173–195 (1998)
10. Hell, P., Nešetřil, J.: On the complexity of H-colouring. Journal of Combinatorial Theory B 48, 92–110 (1990)
11. Kratochvíl, J., Manuel, P., Miller, M.: Generalized domination in chordal graphs. Nordic Journal of Computing 2, 41–50 (1995)
12. Proskurowski, A., Telle, J.A.: Algorithms for vertex partitioning problems on partial k-trees. SIAM J. Discrete Math. 10, 529–550 (1997)
13. Schaefer, T.J.: The complexity of the satisfability problem. In: Proceedings of the 10th Annual ACM Symposium on Theory of Computing, pp. 216–226. ACM Press, New York (1978)
14. Telle, J.A.: Complexity of domination-type problems in graphs. Nordic Journal of Computing 1, 157–171 (1994)
15. Telle, J.A.: Vertex partitioning problems: characterization, complexity and algorithms on partial k-trees, PhD thesis, Department of Computer Science, Universiy of Oregon, Eugene (1994)