Reconfiguration on nowhere dense graph classes

Sebastian Siebertz*
Institute of Informatics, University of Warsaw, Poland
siebertz@mimuw.edu.pl

Submitted: Nov 9, 2017; Accepted: Jul 19, 2018; Published: Aug 10, 2018

Abstract

Let $\mathcal{Q}$ be a vertex subset problem on graphs. In a reconfiguration variant of $\mathcal{Q}$ we are given a graph G and two feasible solutions $S_s, S_t \subseteq V(G)$ of $\mathcal{Q}$ with $|S_s| = |S_t| = k$. The problem is to determine whether there exists a sequence $S_1, \ldots, S_n$ of feasible solutions, where $S_1 = S_s$, $S_n = S_t$, $|S_i| \leqslant k \pm 1$, and each S_{i+1} results from S_i, $1 \leqslant i < n$, by the addition or removal of a single vertex. We prove that for every nowhere dense class of graphs and for every integer $r \geqslant 1$ there exists a polynomial p_r such that the reconfiguration variants of the distance-r independent set problem and the distance-r dominating set problem admit kernels of size $p_r(k)$. If k is equal to the size of a minimum distance-r dominating set, then for any fixed $\varepsilon > 0$ we even obtain a kernel of almost linear size $\mathcal{O}(k^{1+\varepsilon})$. We then prove that if a class $\mathcal{C}$ is somewhere dense and closed under taking subgraphs, then for some value of $r \geqslant 1$ the reconfiguration variants of the above problems on $\mathcal{C}$ are $\mathsf{W}[1]$-hard (and in particular we cannot expect the existence of kernelization algorithms). Hence our results show that the limit of tractability for the reconfiguration variants of the distance-r independent set problem and distance-r dominating set problem on subgraph closed graph classes lies exactly on the boundary between nowhere denseness and somewhere denseness.

Keywords: Reconfiguration; dominating set; independent set; sparse graph classes; nowhere dense graphs

Mathematics Subject Classifications: 05C85, 68R10

*Supported by the National Science Centre of Poland via POLONEZ grant agreement UMO-2015/19/P/ST6/03998, which has received funding from the European Union's Horizon 2020 research and innovation programme (Marie Skłodowska-Curie grant agreement No. 665778).

1 Introduction

In the reconfiguration framework we are not asked to find a feasible solution to an optimization problem $\mathcal{Q}$, but rather to transform a source feasible solution S_s into a more desirable feasible target solution S_t such that each intermediate solution is also feasible. This framework allows to model real-world dynamic situations in which we need to transform one valid system state into another and it is crucial that the system keeps running in all intermediate states.

The literature focuses mainly on the problem of determining the existence of a reconfiguration sequence between two given solutions; an even more difficult problem is to actually find a (possibly minimum-length) reconfiguration sequence of solutions. Typically there are exponentially many feasible solutions to an instance I, and not surprisingly, the above problem has been shown to be PSpace-complete for the reconfiguration variants of many important NP-complete problems.

Reconfiguration problems received considerable attention in recent literature. The studied problems include Vertex Coloring [1, 3, 4, 6, 8, 7], List Edge-Coloring [23], Vertex Cover [33, 35], Independent Set [5, 20, 22, 24, 29], Clique, Set Cover, Integer Programming, Matching, Spanning Tree, Matroid Bases [22], Satisfiability [16, 32, 40], Shortest Path [2, 26], Subset Sum [21], Dominating Set [5, 18, 19, 34], Odd Cycle Transversal, Feedback Vertex Set, and Hitting Set [34]. We refer to the surveys [38, 41] and the thesis [31] for a detailed overview.

A systematic study of the parameterized complexity of reconfiguration problems was initiated by Mouawad et al. [34]. The authors study mostly graph theoretical vertex subset problems, that is, solutions consist of subsets $S \subseteq V(G)$ of the input graph G. For such problems, one natural parameterization is the parameter k, a bound on the size of feasible solutions, another natural parameter is ℓ, the length of the reconfiguration sequence. They proved that Feedback Vertex Set and Bounded Hitting Set (where the cardinality of each set from the input is bounded) admit polynomial reconfiguration kernels (parameterized by k). Concerning lower bounds, they proved that reconfiguration of Dominating Set is $\mathsf{W}[2]$-hard when parameterized by $k+\ell$, as well as a general result on reconfigurations of hereditary properties and their parametric duals, implying $\mathsf{W}[1]$-hardness of reconfiguration of Independent Set, Induced Forest and Bipartite Subgraph parameterized by $k+\ell$, and Vertex Cover, Feedback Vertex Set, and Odd Cycle Transversal parameterized by ℓ.

In this work we consider the token addition and removal (TAR) model of reconfiguration. In this model, for a vertex subset problem $\mathcal{Q}$ on graphs, we are given a graph G and two feasible solutions $S_s, S_t \subseteq V(G)$ of $\mathcal{Q}$ with $|S_s| = |S_t| = k$. The problem is to determine whether there exists a sequence $S_1, \ldots, S_n$ of feasible solutions, where $S_1 = S_s$, $S_n = S_t$, each S_i has size k or $k-1$ if $\mathcal{Q}$ is a maximization problem and size k or $k+1$ if $\mathcal{Q}$ is a minimization problem, and each S_{i+1} results from S_i, $1 \leqslant i < n$, by adding or removing exactly one vertex. Independent Set Reconfiguration is known to be PSpace-complete on graphs of bounded bandwidth [35, 42] and W[1]-hard parameterized by k on general graphs [25]. On the positive side, the problem was shown to be fixed-parameter tractable, with parameter k, for graphs of bounded degree, planar graphs, and graphs excluding $K_{3,d}$ as a subgraph, for any constant d [24, 25]. This result was extended by Lokshtanov et al. [30] to graphs of bounded degeneracy and nowhere dense graphs. Lokshtanov et al. also proved that Dominating Set Reconfiguration is $\mathsf{W}[1]$-hard parameterized by $k+\ell$ on general graphs and fixed-parameter tractable, with parameter k, for graphs excluding $K_{d,d}$ as a subgraph, for any constant d (in particular on degenerate graph classes and nowhere dense classes).

Nowhere dense graph classes, which are also the object of study in the present paper, are very general classes of uniformly sparse graphs [36, 37]. Many familiar classes of sparse graphs, like planar graphs, graphs of bounded treewidth, graphs of bounded degree, and, in fact, all classes that exclude a fixed (topological) minor, are nowhere dense. Notably, classes of bounded average degree or bounded degeneracy are not necessarily nowhere dense. In an algorithmic context this is reasonable, as every graph can be turned into a graph of degeneracy at most 2 by subdividing every edge once; however, the structure of the graph is essentially preserved under this operation. In our context, a particularly interesting algorithmic result states that every first-order definable property of graphs can be decided in almost linear time on nowhere dense graph classes [17]. This result implies that the reconfiguration variants of many of the above mentioned vertex subset problems are fixed-parameter tractable with respect to parameter $k + \ell$ on every nowhere dense graph class (the existence of a reconfiguration sequence can be expressed with $\mathcal{O}(k \cdot \ell)$ quantifiers in first-order logic, whenever the property itself can be defined by a first-order formula), and by the result of [17] can be decided in fixed-parameter time.

Nowhere dense graph classes play a special role for Dominating Set and its more general variant Distance-r Dominating Set. A distance-r dominating set in a graph G (for a fixed integer parameter r) is a set $D \subseteq V(G)$ such that every vertex of G is at distance at most r to a vertex from D. Distance-r Dominating Set was shown to be fixed-parameter tractable on nowhere dense classes in [10] (this result is again implied by the more general result of [17] which was obtained later). It was then shown that nowhere dense classes are the limit of tractability based on sparsity methods, more precisely, it was shown in [13] that if a class $\mathcal{C}$ is not nowhere dense and closed under taking subgraphs, then there is some $r \geqslant 1$ such that Distance-r Dominating Set on $\mathcal{C}$ is $\mathsf{W}[2]$-hard. It was later shown that the problem admits a polynomial kernel [28] and in fact an almost linear kernel [14] on nowhere dense classes.

A kernelization algorithm, or just a kernel, is a polynomial time algorithm which transforms an input instance (G, k) of a parameterized problem to an equivalent instance (G', k') such that $|G'|+k' \leqslant f(k)$ for some function f. Hence for a reconfiguration problem a kernelization algorithm is a polynomial time algorithm which transforms every input instance (G, k, S_s, S_t) into an instance (G', k', S'_s, S'_t) with $|G'| + k' \leqslant f(k)$ for some function f and such that there exists a valid reconfiguration sequence $S_1 = S_s, S_2, \ldots, S_n = S_t$ in G if and only if there exists a valid reconfiguration sequence $S'_1 = S'_s, S'_2, \ldots, S'_m = S'_t$ in G'. Every fixed-parameter tractable problem admits a kernel, however, possibly of exponential or worse size. On the reduced instance (G', k', S'_s, S'_t) one can then run a brute force algorithm to decide whether the initial instance was a positive instance.

1.1 Our results.

We prove that for every nowhere dense class of graphs and for every $r \geqslant 1$ there exists a polynomial p_r such that DISTANCE-r INDEPENDENT SET RECONFIGURATION and DISTANCE-r DOMINATING SET RECONFIGURATION admit kernels of size $p_r(k)$. For DISTANCE-r DOMINATING SET RECONFIGURATION, if k is equal to the size of a minimum distance-r dominating set of G, then for any fixed $\varepsilon > 0$ we even obtain kernels of almost linear size $\mathcal{O}(k^{1+\varepsilon})$.

For DISTANCE-r DOMINATION SET RECONFIGURATION there is a technical subtlety that prevents us from reducing the input instance (G, k, I_s, I_t) to an equivalent instance (G', k, I_s, I_t) such that G' is a subgraph of G. Instead, we can kernelize to an annotated version of the problem, where only a given subset of vertices of G' needs to be dominated, or we can output an instance (G', k, I_s, I_t), where G' does not belong to the class $\mathcal{C}$ under consideration (its density parameters are only slightly larger than those of G, though). Formally, in any case, we do not reduce to the same problem, hence we compute only a so-called *bi-kernel* for the problem. Our results generalize the earlier mentioned results of Lokshtanov et al. [30], who proved that INDEPENDENT SET RECONFIGURATION (i.e. the case $r = 1$) is fixed-parameter tractable on every nowhere dense graph class and DOMINATING SET RECONFIGURATION (i.e. the case $r = 1$) is fixed-parameter tractable if the input graph does not contain large complete bipartite graphs (as a subgraph), in particular on all nowhere dense graph classes.

Our methods for DISTANCE-r INDEPENDENT SET RECONFIGURATION generalize those of Lokshtanov et al. [30] for INDEPENDENT SET RECONFIGURATION to the more general setting of distance-r independence. They are strongly based on the equivalence of nowhere denseness and uniform quasi-wideness (a notion that will be defined in the next section) and polynomial bounds for the quasi-wideness functions which were recently obtained by Kreutzer et al. [28] and Pilipczuk et al. [39].

Our methods for DISTANCE-r DOMINATING SET RECONFIGURATION combine the approach of Lokshtanov et al. [30] for DOMINATING SET RECONFIGURATION with new methods developed for the kernelization of DISTANCE-r DOMINATING SET on nowhere dense graph classes by Eickmeyer et al. [14].

We then prove that if a class $\mathcal{C}$ is somewhere dense and closed under taking subgraphs, then for some value of $r \geqslant 1$ the reconfiguration variants for these problems on $\mathcal{C}$ are $\mathsf{W}[1]$-hard (and in particular we cannot expect the existence of kernelization algorithms). Hence our results show that the limit of tractability for DISTANCE-r INDEPENDENT SET RECONFIGURATION and DISTANCE-r DOMINATING SET RECONFIGURATION on subgraph closed graph classes lies exactly on the boundary between nowhere denseness and somewhere denseness. Our hardness results are rather straightforward generalizations of the $\mathsf{W}[1]$-hardness proofs known for INDEPENDENT SET and DOMINATING SET to their distance-r variants.

2 Preliminaries

Graphs. All graphs in this paper are finite, undirected and simple. Our notation is standard, we refer to the textbook [11] for more background on graphs. We write $V(G)$ for the vertex set of a graph G and $E(G)$ for its edge set. For $r \in \mathbb{N}$, a graph G and $v \in V(G)$ we write $N_r(v)$ for the set of vertices of G at distance at most r from v. The radius of G is the minimum integer r such that there is $v \in V(G)$ with $N_r(v) = V(G)$.

Independent sets and dominating sets. Let G be a graph and $r \in \mathbb{N}$. A set $B \subseteq V(G)$ is called *r-independent* in G if for all distinct $u, v \in B$ we have $\mathrm{dist}_G(u, v) > r$. The set B is a *distance-r dominating set* in G if $N_r(B) = \bigcup_{v \in B} N_r(v) = V(G)$.

Minors and subdivisions. Let G be a graph and let $r \in \mathbb{N}$. A graph H with vertex set $\{v_1, \ldots, v_n\}$ is a *depth-r minor* of G, written $H \preccurlyeq_r G$, if there are connected and pairwise vertex disjoint subgraphs $H_1, \ldots, H_n \subseteq G$, each of radius at most r, such that if $v_i v_j \in E(H)$, then there are $w_i \in V(H_i)$ and $w_j \in V(H_j)$ with $w_i w_j \in E(G)$.

An *r-subdivision* of H is obtained by replacing edges of H by internally vertex disjoint paths of length (exactly) r. We write H_r for the r-subdivision of H.

Nowhere denseness. A class $\mathcal{C}$ of graphs is *nowhere dense* if there exists a function $t \colon \mathbb{N} \to \mathbb{N}$ such that $K_{t(r)} \not\preccurlyeq_r G$ for all $r \in \mathbb{N}$ and for all $G \in \mathcal{C}$. Otherwise, $\mathcal{C}$ is called *somewhere dense*.

Uniform quasi-wideness. A class $\mathcal{C}$ of graphs is called *uniformly quasi-wide* if there are functions $N \colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ and $s : \mathbb{N} \to \mathbb{N}$ such that for all $r, m \in \mathbb{N}$ and all subsets $A \subseteq V(G)$ for $G \in \mathcal{C}$ of size $|A| \geqslant N(r, m)$ there is a set $S \subseteq V(G)$ of size $|S| \leqslant s(r)$ and a set $B \subseteq A \setminus S$ of size $|B| \geqslant m$ which is r-independent in $G - S$.

It was shown by Nešetřil and Ossona de Mendez [37] that a class $\mathcal{C}$ of graphs is nowhere dense if and only if it is uniformly quasi-wide. Quasi-wideness is a very useful property for distance-r domination, as large $2r$-independent sets are natural obstructions for small distance-r dominating sets. For us it will be important that the function N can be assumed to be polynomial in m (the degree of the polynomial may depend on r) and that the sets B and S can be efficiently computed. Polynomial bounds were first obtainend by Kreutzer et al. [28], we refer to the improved bounds of Pilipczuk et al. [39].

Lemma 1 (Pilipczuk et al. [39])**.** *Let $\mathcal{C}$ be a nowhere dense class of graphs. For all $r \in \mathbb{N}$ there is a polynomial $N_r \colon \mathbb{N} \to \mathbb{N}$ and a constant $t_r \in \mathbb{N}$, such that the following holds. Let $G \in \mathcal{C}$ and let $A \subseteq V(G)$ be a vertex subset of size at least $N_r(m)$, for a given m. Then there exists a set $S \subseteq V(G)$ of size $|S| \leqslant t_r$ and a set $B \subseteq A \setminus S$ of size $|B| \geqslant m$ which is r-independent in $G - S$. Moreover, given G and A, such sets S and B can be computed in time $\mathcal{O}(|A| \cdot |E(G)|)$.*

We remark that the $\mathcal{O}$-notation in the above lemma hides constant factors depending on r (which is considered fixed) and the class $\mathcal{C}$.

A-avoiding paths. Let G be a graph and let $A \subseteq V(G)$ be a subset of vertices. For vertices $v \in A$ and $u \in V(G)$, a path P connecting u and v is called *A-avoiding* if all its vertices apart from u and v do not belong to A.

Projection profiles. Let G be a graph, $A \subseteq V(G)$ and $r \in \mathbb{N}$. The *r-projection* of a vertex $u \in V(G)$ onto A, denoted $M_r^G(u, A)$ is the set of all vertices $v \in A$ that can be connected to u by an A-avoiding path of length at most r. The *r-projection profile* of a vertex $u \in V(G)$ on A is the function $\rho_r^G[u, A]$ mapping vertices of A to $\{0, 1, \ldots, r, \infty\}$, defined as follows: for every $v \in A$, the value $\rho_r^G[u, A](v)$ is the length of a shortest A-avoiding path connecting u and v and ∞ in case this length is larger than r. We define

$$\widehat{\mu}_r(G, A) = |\{\rho_r^G[u, A] \colon u \in V(G)\}|$$

to be the number of different r-projection profiles realized on A. For $u, v \in V(G)$ we define

$$u \cong_{A,r} v \quad \iff \quad \rho_r^G[u, A] = \rho_r^G[v, A].$$

Lemma 2 (Eickmeyer et al. [14]). *Let $\mathcal{C}$ be a nowhere dense class of graphs. Then there is a function $f_{\mathrm{proj}}(r, \varepsilon)$ such that for every $r \in \mathbb{N}$, $\varepsilon > 0$, graph $G \in \mathcal{C}$, and vertex subset $A \subseteq V(G)$, it holds that $\widehat{\mu}_r(G, A) \leqslant f_{\mathrm{proj}}(r, \varepsilon) \cdot |A|^{1+\varepsilon}$.*

We remark that in [14] r-projections onto A are defined only for vertices which do not lie inside A themselves. This does not affect the statement of Theorem 2, as this change of definition accounts only for a term $|A|$, which can be absorbed in the function f_{proj}.

Parameterized complexity. A problem is fixed-parameter tractable on a class $\mathcal{C}$ of graphs with respect to a parameter k, if there is an algorithm deciding whether a graph $G \in \mathcal{C}$ admits a solution of size k in time $f(k) \cdot |V(G)|^c$, for a computable function f and constant c. A kernelization algorithm is a polynomial time algorithm which reduces the input instance to a sub-instance of size bounded in the parameter only (independently of the input graph size). Every fixed-parameter tractable problem admits a kernel, however, possibly of exponential or worse size. For efficient algorithms it is therefore most desirable to obtain polynomial, or optimally even linear, kernels. The W-hierarchy is a collection of parameterized complexity classes $\mathsf{FPT} \subseteq \mathsf{W}[1] \subseteq \mathsf{W}[2] \subseteq \ldots$. The assumption $\mathsf{FPT} \subsetneq \mathsf{W}[1]$ can be seen as the analogue of the conjecture that $\mathsf{P} \subsetneq \mathsf{NP}$. Therefore, showing hardness in the parameterized setting is usually accomplished by establishing an *fpt*-reduction to a $\mathsf{W}[1]$-hard problem. We refer to the textbooks [9, 12, 15] for extensive background on parameterized complexity.

Reconfiguration. The token addition and removal variant of the Distance-r Independent Set Reconfiguration problem (r-ISR) is defined as follows. On input (G, k, I_s, I_t), where G is a graph, $k \in \mathbb{N}$, and I_s, I_t are distance-r independent sets in G of size k, the problem is to determine whether there exists a sequence $I_s = I_1, \ldots, I_\ell = I_t$ such that for all $1 \leqslant j \leqslant \ell$

1. I_j is a distance-r independent set in G,
2. $|\Delta(I_j, I_{j+1})| = |(I_j \setminus I_{j+1}) \cup (I_{j+1} \setminus I_j)| = 1$, and
3. $k - 1 \leqslant |I_j| \leqslant k$.

The Distance-r Dominating Set Reconfiguration problem (r-DSR) is defined analogously, we only demand that in the fourth item we have $k \leqslant |D_j| \leqslant k+1$ for the appearing distance-r dominating sets D_j, $1 \leqslant j \leqslant \ell$. We obtain positive results also for the variants where for r-ISR we get as input two integer parameters k, k' and we replace the fourth condition by $k \leqslant |I_j| \leqslant k'$ for all $1 \leqslant j \leqslant \ell$. For r-DSR we may remove the condition on a lower bound completely, that is, in the forth condition demand only that $|D_j| \leqslant k+1$ for all $i \leqslant j \leqslant \ell$.

3 Distance-r independent set reconfiguration

3.1 Polynomial kernel

Our approach for kernelization of r-ISR is similar to that of Lokshtanov et al. [30]. We iteratively remove *irrelevant vertices* from the input instance, until this is no longer possible, in which case the resulting instance will be small.

Irrelevant vertices. Let (G, I_s, I_t, k) be an instance of Distance-r Independent Set Reconfiguration. A vertex $v \in V(G) \setminus (I_s \cup I_t)$ is called an *irrelevant vertex* if (G, k, I_s, I_t) is a positive instance if and only if $(G - v, k, I_s, I_t)$ is a positive instance.

The following lemma shows that given G is large, we can efficiently find an irrelevant vertex.

Lemma 3. *Let (G, k, I_s, I_t) for $G \in \mathcal{C}$ be an instance of* Distance-r Independent Set Reconfiguration, *$M \coloneqq I_s \cup I_t$ and let $R \coloneqq V(G) \setminus M$. Let $S \subseteq V(G)$ and $B \subseteq R \setminus S$ such that B is $2r$-independent in $G - S$. Furthermore, assume that all vertices of B have the same r-projection profile to S, i.e., $\rho^G_r[u, S] = \rho^G_r[v, S]$. If $|B| \geqslant 2k$, then any $v \in B$ is an irrelevant vertex.*

Proof. Let $v \in B$ and enumerate $2k-1$ vertices of $B \setminus \{v\}$ as $w_1, \ldots, w_{2k-1}$. We aim to show that v is an irrelevant vertex. Observe that since $B \subseteq R \setminus S$ the set $\{v, w_1, \ldots, w_{2k-1}\}$ is disjoint from the set $M \cup S$.

Consider a reconfiguration sequence $I_s = I_1, I_2, \ldots, I_t = I_\ell$ from I_s to I_t in G with a minimum number of occurrences of v. We want to prove that v does not occur at all in the sequence, as this proves that v is irrelevant. Towards a contradiction assume that v does occur in the sequence and let p, $1 < p < \ell$, be the first index at which v appears in I_p (that is, $v \in I_p$ and $v \notin I_i$ for all $i < p$). Let $q+1, p < q+1 \leqslant \ell$ be the first index after p at which v is removed (that is, $v \in I_p, \ldots, I_q$ and $v \notin I_{q+1}$). We will modify the sub-sequence $I_p, \ldots, I_q$ such that it does not use v, contradicting our choice of a reconfiguration sequence with a minimum number of occurrences of v. Fix some j, $p \leqslant j \leqslant q$, and let $I = I_j \setminus \{v\}$ and $I' = I_{j+2} \setminus \{v\}$.

Claim 1. *If there is $z \in I$ with $\mathrm{dist}_G(w_i, z) \leqslant r$ for some w_i, then $\mathrm{dist}_G(w_\ell, z) > r$ for all $\ell \neq i$.*

Proof. Assume towards a contradiction that there is $\ell \neq i$ with $\mathrm{dist}_G(w_\ell, z) \leqslant r$. Let P_i be a shortest path (of length at most r) between w_i and z and let P_ℓ a shortest path (of length at most r) between w_ℓ and z. As B is $2r$-independent in $G - S$ there exists a vertex $s \in S$ with $s \in V(P_i)$ or $s \in V(P_\ell)$. By symmetry we may assume that $s \in V(P_i)$ and assume that among all vertices of S which lie on P_i, the vertex s is the one which is closest to w_i. Then we have $\mathrm{dist}_G(w_i, z) = \mathrm{dist}_G(w_i, s) + \mathrm{dist}_G(s, z)$. Now we have $\rho_r^G[v, S] = \rho_r^G[w_i, S]$, hence $\mathrm{dist}_G(v, s) = \mathrm{dist}_G(w_i, s)$. This implies

$$\mathrm{dist}_G(v, z) \leqslant \mathrm{dist}_G(v, s) + \mathrm{dist}_G(s, z) = \mathrm{dist}_G(w_i, s) + \mathrm{dist}_G(s, z) = \mathrm{dist}_G(w_i, z) \leqslant r,$$

contradicting that I_j is a distance-r independent set. ⌟

Claim 2. *There exists $w \in \{w_1, \ldots, w_{2k-1}\}$ with $(I \cup I') \cap N_r(w) = \varnothing$.*

Proof. For $z \in I$, if there is $w_i \in \{w_1, \ldots, w_{2k-1}\}$ with $z \in N_r(w_i)$, i.e., $\mathrm{dist}_G(z, w_i) \leqslant r$, then by Theorem 1 we have $\mathrm{dist}_G(w_\ell, z) > r$ for all $\ell \neq i$. As the set I contains only $k-1$ elements we conclude that there are k elements $u_1, \ldots, u_k$ in $\{w_1, \ldots, w_{2k-1}\}$ with $I \cap N_r(u_i) = \varnothing$ for all $1 \leqslant i \leqslant k$. We apply the same reasoning to the set I' and $\{u_1, \ldots, u_k\}$ (note that the claims are also applicable to I', as j is chosen arbitrary). This leaves us with an element $w \in \{w_1, \ldots, w_{2k-1}\}$ with $(I \cup I') \cap N_r(w) = \varnothing$. ⌟

As j was chosen arbitrary, we conclude that for every j there exists an element $w^j \in \{w_1, \ldots, w_{2k-1}\}$ such that $(I_j \setminus \{v\}) \cup \{w^j\}$ and $(I_{j+2} \setminus \{v\}) \cup \{w^j\}$ are distance-r independent sets of the same size as I_j. Note that we have $|I_p| = |I_q| = k$, as v was introduced at I_p and removed at I_{q+1}. Hence, if we have $j = p + 2x$ for some $x \in \mathbb{N}$, then $I_{j+1} \subseteq I_j, I_{j+2}$, hence also $(I_{j+1} \setminus \{v\}) \cup \{w^j\}$ is a distance-r independent sets of size $k-1$.

We now modify the sequence $I_p, \ldots, I_q$ as follows. For each $j = p + 2x$ for some $x \in \mathbb{N}$ such that $p \leqslant j < q$, we replace the subsequence $I_j \to I_{j+1}$ of the reconfiguration sequence by the sequence

$$(I_j \setminus \{v\}) \cup \{w^j\} \to (I_{j+1} \setminus \{v\}) \cup \{w^j\} \to (I_{j+2} \setminus \{v\}) \cup \{w^j\} \to (I_{j+2} \setminus \{v\})$$

and we replace I_q by $(I_q \setminus \{v\}) \cup \{w^q\}$.

By our above argument, each of the intermediate configurations is a distance-r independent set of size k or $k-1$. Furthermore, the transition I_{p-1}, J_p is valid, so is every intermediate transition and the transition J_q, I_{q+1}. This finishes the proof of the lemma. □

Theorem 4. *Let $\mathcal{C}$ be a nowhere dense class of graphs and let $r \in \mathbb{N}$. Let (G, k, I_s, I_t) be an instance of* Distance-r Independent Set Reconfiguration*, where $G \in \mathcal{C}$. Then we can compute in polynomial time a subgraph $G' \subseteq G$ with $I_s, I_k \subseteq V(G')$ such that (G, k, I_s, I_t) is a positive instance if and only if (G', k, I_s, I_t) is a positive instance and G' has order polynomial in k.*

Proof. Let $N = N_{2r} : \mathbb{N} \to \mathbb{N}$ be the function and $t = t_{2r} \in \mathbb{N}$ be the constant describing $\mathcal{C}$ as uniformly quasi-wide (for parameter $2r$) as defined in Theorem 1. Let $M \coloneqq I_s \cup I_t$ of size $2k$ and let $R \coloneqq V(G) \setminus M$.

If $|R| \geqslant N(2k \cdot (r+2)^t)$, according to Theorem 1 we can compute in polynomial time a set $S \subseteq V(G)$ of size $|S| \leqslant t$ and a set $B \subseteq R \setminus S$ of size $|B| \geqslant 2k(r+2)^t$ which is $2r$-independent in $G - S$. We classify the elements of B with respect to their r-projections to the set S. The corresponding equivalence relation $\cong_{S,r}$ on B has at most $(r+2)^t$ equivalence classes, as $|S| \leqslant t$ and $\rho_r^G[u, S]$ is a mapping from S to $\{0, 1, \ldots, r, \infty\}$. Since $|B| \geqslant 2k(r+2)^t$, we know that at least one equivalence class contains at least $2k$ vertices of B. We apply Theorem 3 to this equivalence class to find an irrelevant vertex v. We remove v from the graph and iterate this procedure until $|R| < N(2k \cdot (r+2)^t)$. In this case the resulting graph has size at most $N(2k \cdot (r+2)^t) + 2k$, which is polynomial in k for each fixed value of r. □

It is easy to see that we can carry out the same proof for the reconfiguration variant where we get as input two integer parameters k, k' and we replace the fourth condition by $k \leqslant |I_j| \leqslant k'$ for all $1 \leqslant j \leqslant \ell$. The kernel will have size polynomial in k'.

We remark that the kernel does possibly not preserve the length of a shortest reconfiguration sequence. It remains an interesting open question to compute a kernel with this preservation property.

3.2 Lower bounds

Recall that for a graph G and $s \in \mathbb{N}$, G_s denotes the s-subdivision of G. Our hardness result is based on the following observation by Nešetřil and Ossona de Mendez.

Lemma 5 (Nešetřil and Ossona de Mendez [37], see also [13])**.** *Let $\mathcal{C}$ be somewhere dense and closed under taking subgraphs. Then there is $s \in \mathbb{N}$ such that for all graphs G we have $G_s \in \mathcal{C}$.*

Furthermore, we use that Independent Set Reconfiguration, i.e., the case $r = 1$ is hard.

Lemma 6 (Ito et al. [25])**.** Independent Set Reconfiguration *is* $\mathsf{W}[1]$*-hard.*

Theorem 7. *Let $\mathcal{C}$ be somewhere dense and closed under taking subgraphs. Then there is $r \in \mathbb{N}$ such that* Distance-r Independent Set Reconfiguration *is* $\mathsf{W}[1]$*-hard on $\mathcal{C}$.*

Proof. According to Theorem 5, there is $s \in \mathbb{N}$ such that for all graphs G we have $G_s \in \mathcal{C}$. We reduce 1-ISR to $(4s-1)$-ISR on $\mathcal{C}$ by establishing the following. For each graph G there exists a polynomial time computable graph $H \in \mathcal{C}$ such that $V(G) \subseteq V(H)$ and such that

1. every independent set I in G is a $(4s-1)$-independent set in H and
2. every $(4s-1)$-independent set I of size at least 2 in H consists only of vertices which are also vertices of G and I is an independent set in G.

Note that we may assume that the parameter k is always at least 2. The above properties guarantee that every reconfiguration sequence $I_1, \ldots, I_\ell$ of independent sets in G corresponds uniquely to a reconfiguration sequence of distance-$(4s-1)$ independent sets in H and vice versa. Hence, we can conclude the statement of the theorem for $r = 4s-1$ by applying Theorem 6. Note that the reduction also establishes $\mathsf{W}[1]$-hardness of Distance-r Independent Set on somewhere dense graph classes which are closed under taking subgraphs.

Let G be a graph with vertex set $\{v_1, \ldots, v_n\}$ and edge set $\{e_1, \ldots, e_m\}$. We remark that the hardness result of Theorem 6 works also if we assume that all input graphs do not have isolated vertices, so we may assume that G does not contain isolated vertices. We define the new graph J with vertex set

$$\{v_1, \ldots, v_n, e_1, \ldots, e_m, w\}$$

and edge set

$$\{ve : v \in V(G), e \in E(G), v \in e\} \cup \{ew : e \in E(G)\}.$$

We claim that $H = J_s$ satisfies the above claimed properties.

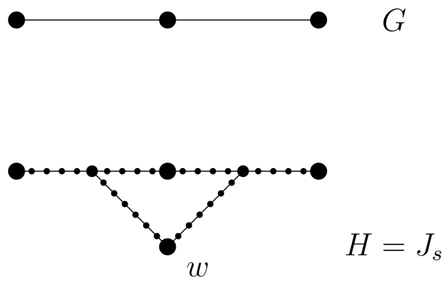

Figure 1: A graph G and the constructed graph $H \in \mathcal{C}$. Vertices at distance 1 in G have distance $2s$ in H, while vertices at distance 2 have distance $4s$. All vertices of H which do not correspond to vertices of G have distance at most $4s-1$.

Let I be a distance-1 independent set in G. By construction of J, if $u, v \in V(G)$ are adjacent in G, then they have distance $2s$ in J_s, otherwise they have distance $4s$ in J_s (via the vertex w). Hence, I is a distance-$(4s-1)$ independent set in J_s.

Conversely, let I be a distance-$(4s-1)$ independent set in J_s of size at least 2. First observe that I consists only of vertices which are also vertices of G. All other vertices have mutual distance at most $4s-1$ via the vertex w. As seen above, the elements of I have distance $4s$ in J_s and therefore distance at least 2 in G, that is, I is an independent set in G. This finishes the proof. $\square$

4 Distance-r dominating set reconfiguration

4.1 Polynomial kernel

The kernelization for Distance-r Dominating Set Reconfiguration strongly depends on the following notion of a domination core which was (in different variants) also used in the earlier kernelization results for distance-r dominating sets [10, 13, 14].

Domination core. Let G be a graph and $k, r \in \mathbb{N}$. A set $Z \subseteq V(G)$ is called a (k, r)-*domination core* if every set D of size at most k that r-dominates Z also r-dominates G.

Clearly, $V(G)$ is a (k, r)-domination core. Hence, starting with $Z = V(G)$, using the next lemma, we can gradually remove vertices from Z while maintaining the invariant that Z is a (k, r)-domination core. The proof of the lemma is the same as the proof of Lemma 11 in [10] and Lemma 4.1 in [28], we just use the better bounds from Theorem 1.

Lemma 8. *Let $\mathcal{C}$ be a nowhere dense class of graphs and let $k, r \in \mathbb{N}$. Let $N = N_{2r}$ and $t = t_{2r}$ be the functions characterizing $\mathcal{C}$ as uniformly quasi-wide according to Theorem 1 with parameter $2r$. There is an algorithm that, given a graph $G \in \mathcal{C}$, $k \in \mathbb{N}$ and $Z \subseteq V(G)$ with $|Z| > N\big((k+2)(2r+1)^t\big) =: \ell$ runs in time $\mathcal{O}(t \cdot \ell \cdot |E(G)|)$, and returns a vertex $w \in Z$ such that for any set $X \subseteq V(G)$ with $|X| \leqslant k$, it holds that X is an r-dominating set of Z if, and only if, X is an r-dominating set of $Z \setminus \{w\}$.*

We iteratively apply Theorem 8 for at most n times, until this is no longer possible. This yields the following lemma.

Lemma 9. *Let $\mathcal{C}$ be a nowhere dense class of graphs and let $k, r \in \mathbb{N}$. There exists a polynomial q_r and a polynomial time algorithm that, given a graph $G \in \mathcal{C}$ and $k \in \mathbb{N}$ either correctly concludes that G cannot be r-dominated by a set of at most k vertices, or finds a (k, r)-domination core $Z \subseteq V(G)$ of G of size at most $q_r(k)$.*

We now define the annotated problem Distance-r Z-Dominating Set, r-ZDS, as the problem to find on input (G, Z, k) a set D with $Z \subseteq N_r(D)$. Such a set D is called a (Z, r)-*dominator*. By definition, if Z is a (k, r)-domination core, then every (Z, r)-dominator of size at most k corresponds to a distance-r dominating set of G. On the other hand, every distance-r dominating set of G in particular dominates every subset $Z \subseteq V(G)$. We define the reconfiguration variant of the problem, r-ZDSR, in the obvious way.

Theorem 10. *Let $\mathcal{C}$ be a nowhere dense class of graphs and let $r \in \mathbb{N}$. Let (G, k, D_s, D_t) be an instance of r-DSR, where $G \in \mathcal{C}$. We can compute in polynomial time a subgraph $G' \subseteq G$ with $D_s, D_k \subseteq G'$ and $Z \subseteq V(G')$ such that (G, k, D_s, D_t) is a positive instance of r-DSR if and only if $((G', Z), k, D_s, D_t)$ is a positive instance of r-ZDSR and G' has order polynomial in k.*

Proof. We compute a (k, r)-domination core $Z \subseteq V(G)$ of size at most $q_r(k)$ using Theorem 9. Let $\varepsilon > 0$ and let $f_{\text{proj}}(r, \varepsilon)$ be the function from Theorem 2. According to the lemma there are at most $f_{\text{proj}}(r,\varepsilon)\cdot|Z|^{1+\varepsilon}$ different r-projections to Z. Recall that $u\cong_{Z,r} v \Leftrightarrow \rho_r^G[u,Z]=\rho_r^G[v,Z]$. Now for each projection class κ we choose a representative v_κ from that class.

Claim 1. *For all $u,v\in V(G)$, $\rho_r^G[u,Z]=\rho_r^G[v,Z]$ implies $N_r^G(u)\cap Z=N_r^G(v)\cap Z$.*

Proof. Let $z\in N_r^G(u)\cap Z$ and let P be a shortest path between u and z. If P is Z-avoiding we conclude from $\rho_r^G[u,Z]=\rho_r^G[v,Z]$ that there exists also a Z-avoiding path P' of the same length as P between v and z, which implies $z\in N_r^G(v)\cap Z$. Otherwise, let z' be the vertex of Z on P which is closest to u and let Q be the initial part of P between u and z'. Note that Q is a shortest path between u and z'. By the same argument as above, we find a Z-avoiding path Q' of the same length as Q between v and z'. By replacing Q in P by Q' we obtain a path of the same length at P between v and z, which again proves $z\in N_r^G(v)\cap Z$. ⌟

We now construct G' such that it contains D_s, D_t, the set Z, all the representatives v_κ and furthermore a small set T of vertices such that $N_r^G(v_\kappa)\cap Z=N_r^{G'}(v_\kappa)\cap Z$. The set T is constructed as follows. For each $v\in V(G)$, let T_v be a breadth-first search tree with root v of depth r which has elements of Z as its leaves. Clearly, $\text{dist}_G(v,z)=\text{dist}_{T_v}(v,z)$ for all $z\in N_r(v)\cap Z$. Let T be the set $\bigcup_\kappa V(T_{v_\kappa})\cup\bigcup_{v\in D_s\cup D_t}V(T_v)$. Hence for each v_κ we have $N_r^{G'}(v_\kappa)\cap Z=N_r^G(v_\kappa)\cap Z$, which immediately implies the next claim.

Claim 2. *Let D' be a (Z,r)-dominator in G' which contains only representative vertices v_κ. Then D' is also a (Z,r)-dominator in G.* ⌟

We will always find (Z,r)-dominators of the above form.

Claim 3. *Let $v\in V(G')$. Then there is $v_\kappa\in G'$ such that $N_r^{G'}(v)\cap Z\subseteq N_r^{G'}(v_\kappa)\cap Z$.*

Proof. Let κ be the equivalence class of v in the relation $\cong_{Z,r}$. Then $N_r^{G'}(v)\cap Z\subseteq N_r^{G'}(v_\kappa)\cap Z$. ⌟

Conversely, (Z,r)-dominators in G can be translated to (Z,r)-dominators in G'.

Claim 4. *Let D be a distance-r dominating set in G. Then $D'=\{v_\kappa : v\in D, v\cong_{Z,r} v_\kappa\}$ is a (Z,r)-dominator in G'.*

Proof. As v_κ is chosen so that $\rho_r^G[v,Z]=\rho_r^G[v_\kappa,Z]$, by Claim 1 it holds that $N_r^G(v)\cap Z=N_r^G(v_\kappa)\cap Z$. Hence $Z\subseteq N_r^G(D)$ and $N_r^{G'}(v_\kappa)\cap Z=N_r^G(v_\kappa)\cap Z$ implies that also $Z\subseteq N_r^{G'}(D')$. ⌟

We can now prove that the instance $((G',Z),k,D_s,D_t)$ of r-ZDSR is equivalent to the instance (G,k,D_s,D_t). If $D_1,\ldots,D_n$ is a valid reconfiguration sequence in G, then according to Claim 4 the corresponding sequence $D_1',\ldots,D_n'$ is also a valid reconfiguration sequence of (Z,r)-dominators in G'.

Conversely, Let $D'_1, \ldots, D'_n$ be a reconfiguration sequence of (Z, r)-dominators in G'. We first modify D'_i such that it uses only representative vertices v_κ, using Claim 3. Now according to Claim 2, D'_i is also a (Z, r)-dominator in G. By definition of a (k, r)-domination core, D'_i is a distance-r dominating set in G.

It remains to estimate the size of G'. According to Theorem 9, Z has polynomial size at most $q_r(k)$. According to Theorem 2 there are at most $f_{\text{proj}}(r, \varepsilon) \cdot |Z|^{1+\varepsilon}$ projection classes, hence we add at most so many vertices v_κ to G'. Furthermore, each spanning tree T_{v_κ} has order at most $r \cdot |Z|$. Together with the $2k$ spanning trees we add for D_s and D_t, we have $|V(G')| \leqslant (f_{\text{proj}}(r, \varepsilon) + 2k) \cdot q_r(k)^{2+\varepsilon} \cdot r$, which is polynomial for every fixed value of r and ε. □

The annoying fact that we reduce to an annotated version of the problem can be dealt with by introducing a simple gadget to G'. The same problem occurred also in the kernelization algorithms for DISTANCE-r DOMINATING SET on bounded expansion and nowhere dense graph classes [13, 14]. We refer to these papers for the (very simple) details.

We can find much smaller domination cores if we make a further assumption on the dominating set size. The following definition was first given in [13] and is also the basis for the kernelization of DISTANCE-r DOMINATING SET in [14].

Minimum domination core. Let G be a graph and $r \geqslant 1$. A set $Z \subseteq V(G)$ is a *distance-r domination core* if every set D of minimum size that r-dominates Z also r-dominates G.

The little change in the definition makes a large difference for the sizes of the respective cores, as the next lemma shows.

Lemma 11 (Eickmeyer et al. [14])**.** *Let $\mathcal{C}$ be a nowhere dense class of graphs. There exists a function $f_{\text{core}}(r, \varepsilon)$ and a polynomial-time algorithm that, given a graph $G \in \mathcal{C}$, integer $k \in \mathbb{N}$ and $\varepsilon > 0$, either correctly concludes that G cannot be r-dominated by k vertices, or finds a distance-r domination core $Z \subseteq V(G)$ of G of size at most $f_{\text{core}}(r, \varepsilon) \cdot k^{1+\varepsilon}$.*

If we make the assumption that the source and target sets D_s and D_t are of minimum size, we can work with the improved bounds of Theorem 11 instead of the polynomial bounds of Theorem 9. The final obstacle is to better control the sizes of the trees T_{v_κ} that we add to the graph G' in the construction of Theorem 10. For this, we need the following two lemmas.

Lemma 12 (Lemma 2.9 of [13], adjusted (see Lemma 8 of [14]))**.** *There exists a function $f_{\text{cl}}(r, \varepsilon)$ and a polynomial-time algorithm that, given $G \in \mathcal{C}$, $X \subseteq V(G)$, $r \in \mathbb{N}$, and $\varepsilon > 0$, computes the r-*closure *of X, denoted* $\text{cl}_r(X)$ *with the following properties.*

- *$X \subseteq \text{cl}_r(X) \subseteq V(G)$;*
- *$|\text{cl}_r(X)| \leqslant f_{\text{cl}}(r, \varepsilon) \cdot |X|^{1+\varepsilon}$; and*
- *$|M_r^G(u, \text{cl}_r(X))| \leqslant f_{\text{cl}}(r, \varepsilon) \cdot |X|^{\varepsilon}$ for each $u \in V(G) \setminus \text{cl}_r(X)$.*

We can now compute a breadth-first search tree with root v_κ of depth at most r which stops whenever it first encounters a vertex of $\mathrm{cl}_r(Z)$. This gives us a tree T_{v_κ} of size at most $f_{\mathrm{cl}}(r,\varepsilon)\cdot|Z|^\varepsilon\cdot r$. However, as the breadth-first search does not continue when meeting $\mathrm{cl}_r(Z)$, we now have to connect the vertices of $\mathrm{cl}_r(Z)$ with minimum length paths (up to length r) to preserve all distances. This is possible as the next lemma shows.

Lemma 13 (Lemma 2.11 of [13], adjusted (see Lemma 9 of [14]))**.** *There is a function $f_{\mathrm{pth}}(r,\varepsilon)$ and a polynomial-time algorithm which on input $G\in\mathcal{C}$, $X\subseteq V(G)$, $r\in\mathbb{N}$, and $\varepsilon>0$, computes a superset $X'\supseteq X$ of vertices with the following properties:*

- *whenever* $\mathrm{dist}_G(u,v)\leqslant r$ *for* $u,v\in X$*, then* $\mathrm{dist}_{G[X']}(u,v)=\mathrm{dist}_G(u,v)$*; and*
- $|X'|\leqslant f_{\mathrm{pth}}(r,\varepsilon)\cdot|X|^{1+\varepsilon}$.

We can now prove the following theorem.

Theorem 14. *Let $\mathcal{C}$ be a nowhere dense class of graphs and let $r\in\mathbb{N}$. Let (G,D_s,D_t,k) be an instance of r-DSR, where $G\in\mathcal{C}$ and where D_s and D_t are minimum distance-r dominating sets in G. There is a function $f_{\mathrm{ker}}(r,\varepsilon)$ and a polynomial-time algorithm which on input (G,D_s,D_t,k) computes a subgraph $G'\subseteq G$ with $D_s,D_k\subseteq G'$ and $Z\subseteq V(G')$ such that (G,k,D_s,D_t) is a positive instance of r-DSR if and only if $((G',Z),k,D_s,D_t)$ is a positive instance of r-ZDSR and G' has order at most $f_{\mathrm{ker}}(r,\varepsilon)\cdot k^{1+\varepsilon}$.*

Proof. The proof parallels that of Theorem 10. We compute a distance-r domination core $Z\subseteq V(G)$ using Theorem 11 instead of Theorem 9 (with parameter ε' which will be determined in the course of the proof). Now, using Theorem 12, we compute $Z'=\mathrm{cl}_r(Z)$ and using Theorem 13 we compute $Z''\supseteq Z'$ such that whenever $\mathrm{dist}_G(u,v)\leqslant r$ for $u,v\in Z'$, then $\mathrm{dist}_{G[Z'']}(u,v)=\mathrm{dist}_G(u,v)$. We classify the elements of $V(G)\setminus Z'$ according to their r-projections to Z', that is, we define $u\cong_{Z',r}v\Leftrightarrow\rho_r^G[u,Z']=\rho_r^G[v,Z']$.

We now construct G' such that it contains D_s,D_t, the set Z'', all the representatives v_κ and furthermore a small set T of vertices such that $N_r^G(v_\kappa)\cap Z'=N_r^{G'}(v_\kappa)\cap Z'$. The set T is constructed as follows. For each $v\in V(G)$, let T_v be a breadth-first search tree with root v of depth r which does not continue when meeting Z for the first time. The crucial claim is the following.

Claim 1. *Let v_κ be a representative vertex. Then $\mathrm{dist}_G(v_\kappa,z)=\mathrm{dist}_{G'}(v_\kappa,z)$ for all $z\in N_r^G(v_\kappa)\cap Z$.*

Proof. Let $z\in N_r^G(v_\kappa)\cap Z$ and let P be a minimum length path between v_κ and z. Let z' be the first vertex on P which belongs to Z'. Then we have $\mathrm{dist}_G(v_\kappa,z')=\mathrm{dist}_{T_{v_\kappa}}(v_\kappa,z')$. Now by construction of Z'' we have $\mathrm{dist}_G(z',z)=\mathrm{dist}_{G'}(z',z)$, which implies the claim. ⌟

The rest of the proof works exactly as the proof of Theorem 10. Let us determine a bound on the size of G', which also determines our initial choice of ε'. The set Z has size at most $f_{\mathrm{core}}(r,\varepsilon')\cdot k^{1+\varepsilon'}$. According to Theorem 2 there are at most $f_{\mathrm{proj}}(r,\varepsilon')\cdot|Z|^{1+\varepsilon'}$ projection classes, hence we add at most so many vertices v_κ to G'. The set Z' has size at most $f_{\text{proj}}(r,\varepsilon') \cdot |Z|^{1+\varepsilon'}$ according to Theorem 12 and the set Z'' has size at most $f_{\text{pth}}(r,\varepsilon') \cdot |Z'|^{1+\varepsilon}$ according to Theorem 13. Now, each tree T_v has order at most $|Z'|^{\varepsilon'} \cdot r$. Hence in total we have

$$\begin{aligned} |V(G')| &\leqslant |Z''| + (f_{\text{proj}}(r,\varepsilon') \cdot |Z|^{1+\varepsilon'} + 2) \cdot |Z'|^{\varepsilon'} \cdot r \\ &=: f_{\text{ker}}(r,\varepsilon) \cdot k^{1+\varepsilon} \end{aligned}$$

for an appropriately chosen function f_{ker} and ε'. □

Observe that the constructed kernel preserves shortest reconfiguration sequences.

4.2 Lower bounds

Theorem 15. *Let $\mathcal{C}$ be somewhere dense and closed under taking subgraphs. Then there is $r \in \mathbb{N}$ such that* Distance-r Dominating Set Reconfiguration *is* W[2]*-hard on $\mathcal{C}$.*

Proof. The proof works in principle as the proof of Theorem 7. Again, let $s \in \mathbb{N}$ be the number such that according to Theorem 5 for all graphs G we have $G_s \in \mathcal{C}$. We reduce 1-DSR to $(3s)$-DSR on $\mathcal{C}$ by finding an appropriate subdivision of a graph in which dominating sets are translated 1-to-1 to distance-$(3s)$ dominating sets. Here we can directly use the reduction from set cover to distance-r dominating set from [13], where we use the fact that dominating set and set cover are equivalent problems (just define the set system consisting of the neighborhoods of all vertices). Now use that the reconfiguration variant of dominating set is W[2]-hard [34]. □

5 Conclusion

The study of computationally hard problems on restricted classes of inputs is a very fruitful line of research in algorithmic graph structure theory and in particular in parameterized complexity theory. This research is based on the observation that many problems such as Dominating Set, which are considered intractable in general, can be solved efficiently on restricted graph classes. Of course it is a very desirable goal in this line of research to identify the most general classes of graphs on which certain problems can be solved efficiently. In this work we were able to identify the exact limit of tractability for the reconfiguration variants of the distance-r independent set problem and distance-r dominating set problem on subgraph closed graph classes. Clearly, the main open question is to identify the most general graph classes which are not subgraph closed on which these problems admit efficient solutions.

References

[1] Marthe Bonamy and Nicolas Bousquet. Recoloring bounded treewidth graphs. *Electronic Notes in Discrete Mathematics*, 44:257–262, 2013.

[2] Paul Bonsma. The complexity of rerouting shortest paths. *Theoretical computer science*, 510:1–12, 2013.

[3] Paul Bonsma and Luis Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoretical Computer Science*, 410(50):5215–5226, 2009.

[4] Paul Bonsma, Amer E Mouawad, Naomi Nishimura, and Venkatesh Raman. The complexity of bounded length graph recoloring and csp reconfiguration. In *International Symposium on Parameterized and Exact Computation*, pages 110–121. Springer, 2014.

[5] Nicolas Bousquet, Arnaud Mary, and Aline Parreau. Token jumping in minor-closed classes. *21st International Symposium on Fundamentals of Computation Theory, FCT*, pages 136–149, 2017.

[6] Luis Cereceda, Jan Van Den Heuvel, and Matthew Johnson. Connectedness of the graph of vertex-colourings. *Discrete Mathematics*, 308(5):913–919, 2008.

[7] Luis Cereceda, Jan Van den Heuvel, and Matthew Johnson. Mixing 3-colourings in bipartite graphs. *European Journal of Combinatorics*, 30(7):1593–1606, 2009.

[8] Luis Cereceda, Jan Van Den Heuvel, and Matthew Johnson. Finding paths between 3-colorings. *Journal of graph theory*, 67(1):69–82, 2011.

[9] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, Volume 3. Springer, 2015.

[10] Anuj Dawar and Stephan Kreutzer. Domination problems in nowhere-dense classes. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS* 2009, pages 157–168. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2009.

[11] Reinhard Diestel. *Graph Theory, 4th Edition*, Volume 173 of *Graduate Texts in Mathematics*. Springer, 2012.

[12] Rodney G Downey and Michael R Fellows. *Fundamentals of parameterized complexity*, Volume 4. Springer, 2013.

[13] Pål Grønås Drange, Markus Sortland Dregi, Fedor V. Fomin, Stephan Kreutzer, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, Felix Reidl, Fernando Sánchez Villaamil, Saket Saurabh, Sebastian Siebertz, and Somnath Sikdar. Kernelization and sparseness: the case of Dominating Set. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS*, pages 31:1–31:14. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2016.

[14] Kord Eickmeyer, Archontia C. Giannopoulou, Stephan Kreutzer, O-joung Kwon, Michal Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. Neighborhood complexity and kernelization for nowhere dense classes of graphs. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017*, pages 63:1–63:14, 2017.

[15] Jörg Flum and Martin Grohe. *Parameterized complexity theory.* Springer Science & Business Media, 2006.

[16] Parikshit Gopalan, Phokion G Kolaitis, Elitza Maneva, and Christos H Papadimitriou. The connectivity of boolean satisfiability: computational and structural dichotomies. *SIAM Journal on Computing*, 38(6):2330–2355, 2009.

[17] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. In *Journal of the ACM*, 64(3):17:1–17:32, 2017.

[18] Ruth Haas and Karen Seyffarth. The k-dominating graph. *Graphs and Combinatorics*, 30(3):609–617, 2014.

[19] Arash Haddadan, Takehiro Ito, Amer E Mouawad, Naomi Nishimura, Hirotaka Ono, Akira Suzuki, and Youcef Tebbal. The complexity of dominating set reconfiguration. In *Workshop on Algorithms and Data Structures*, pages 398–409. Springer, 2015.

[20] Robert A Hearn and Erik D Demaine. Pspace-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theoretical Computer Science*, 343(1-2):72–96, 2005.

[21] Takehiro Ito and Erik D Demaine. Approximability of the subset sum reconfiguration problem. In *International Conference on Theory and Applications of Models of Computation*, pages 58–69. Springer, 2011.

[22] Takehiro Ito, Erik D Demaine, Nicholas JA Harvey, Christos H Papadimitriou, Martha Sideri, Ryuhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12-14):1054–1065, 2011.

[23] Takehiro Ito, Marcin Kamiński, and Erik D Demaine. Reconfiguration of list edge-colorings in a graph. *Discrete Applied Mathematics*, 160(15):2199–2207, 2012.

[24] Takehiro Ito, Marcin Kamiński, and Hirotaka Ono. Fixed-parameter tractability of token jumping on planar graphs. In *International Symposium on Algorithms and Computation*, pages 208–219. Springer, 2014.

[25] Takehiro Ito, Marcin Kamiński, Hirotaka Ono, Akira Suzuki, Ryuhei Uehara, and Katsuhisa Yamanaka. On the parameterized complexity for token jumping on graphs. In *International Conference on Theory and Applications of Models of Computation*, pages 341–351. Springer, 2014.

[26] Marcin Kamiński, Paul Medvedev, and Martin Milanič. Shortest paths between shortest paths. *Theoretical Computer Science*, 412(39):5205–5210, 2011.

[27] Marcin Kamiński, Paul Medvedev, and Martin Milanič. Complexity of independent set reconfigurability problems. *Theoretical computer science*, 439:9–15, 2012.

[28] Stephan Kreutzer, Roman Rabinovich, and Sebastian Siebertz. Polynomial kernels and wideness properties of nowhere dense graph classes. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 1533–1545. SIAM, 2017.

[29] Daniel Lokshtanov and Amer E Mouawad. The complexity of independent set reconfiguration on bipartite graphs. *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 185–195, SIAM, 2018.

[30] Daniel Lokshtanov, Amer E Mouawad, Fahad Panolan, MS Ramanujan, and Saket Saurabh. Reconfiguration on sparse graphs. In *Workshop on Algorithms and Data Structures*, pages 506–517. Springer, 2015.

[31] Amer Mouawad. On reconfiguration problems: Structure and tractability. 2015.

[32] Amer E Mouawad, Naomi Nishimura, Vinayak Pathak, and Venkatesh Raman. Shortest reconfiguration paths in the solution space of boolean formulas. In *International Colloquium on Automata, Languages, and Programming*, pages 985–996. Springer, 2015.

[33] Amer E Mouawad, Naomi Nishimura, and Venkatesh Raman. Vertex cover reconfiguration and beyond. In *International Symposium on Algorithms and Computation*, pages 452–463. Springer, 2014.

[34] Amer E Mouawad, Naomi Nishimura, Venkatesh Raman, Narges Simjour, and Akira Suzuki. On the parameterized complexity of reconfiguration problems. *Algorithmica*, 78(1):274–297, 2017.

[35] Amer E Mouawad, Naomi Nishimura, Venkatesh Raman, and Marcin Wrochna. Reconfiguration over tree decompositions. In *International Symposium on Parameterized and Exact Computation*, pages 246–257. Springer, 2014.

[36] Jaroslav Nešetřil and Patrice Ossona de Mendez. First order properties on nowhere dense structures. *The Journal of Symbolic Logic*, 75(03):868–887, 2010.

[37] Jaroslav Nešetřil and Patrice Ossona de Mendez. On nowhere dense graphs. *European Journal of Combinatorics*, 32(4):600–617, 2011.

[38] Naomi Nishimura. Introduction to reconfiguration. *Algorithms 11(4):52*, 2018.

[39] Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. On the number of types in sparse graphs. *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018*, pages 799-808, ACM, 2017.

[40] Konrad W Schwerdtfeger. A computational trichotomy for connectivity of boolean satisfiability. *JSAT*, 8(3/4):173–195, 2014.

[41] Jan van den Heuvel. The complexity of change. *Surveys in combinatorics*, 409:127–160, 2013.

[42] Marcin Wrochna. Reconfiguration in bounded bandwidth and treedepth. arXiv:1405.0847, 2014.