On the complexity of finding large odd induced subgraphs and odd colorings

Rémy Belmonte, Ignasi Sau

On the complexity of finding large odd induced subgraphs and odd colorings*

Rémy Belmonte[1] and Ignasi Sau[2]

[1] University of Electro-Communications, Chofu, Japan
remy.belmonte@gmail.com
[2] LIRMM, Université de Montpellier, CNRS, Montpellier, France
ignasi.sau@lirmm.fr

Abstract. We study the complexity of the problems of finding, given a graph G, a largest induced subgraph of G with all degrees odd (called an *odd* subgraph), and the smallest number of odd subgraphs that partition $V(G)$. We call these parameters $\mathsf{mos}(G)$ and $\chi_{\mathsf{odd}}(G)$, respectively. We prove that deciding whether $\chi_{\mathsf{odd}}(G) \leq q$ is polynomial-time solvable if $q \leq 2$, and NP-complete otherwise. We provide algorithms in time $2^{\mathcal{O}(\mathsf{rw})} \cdot n^{\mathcal{O}(1)}$ and $2^{\mathcal{O}(q \cdot \mathsf{rw})} \cdot n^{\mathcal{O}(1)}$ to compute $\mathsf{mos}(G)$ and to decide whether $\chi_{\mathsf{odd}}(G) \leq q$ on n-vertex graphs of rank-width at most rw, respectively, and we prove that the dependency on rank-width is asymptotically optimal under the ETH. Finally, we give some tight bounds for these parameters on restricted graph classes or in relation to other parameters.

Keywords: odd subgraph; odd coloring; rank-width; parameterized complexity; single-exponential algorithm; Exponential Time Hypothesis.

1 Introduction

Gallai proved, around 60 years ago, that the vertex set of every graph can be partitioned (in polynomial time) into two sets, each of them inducing a subgraph in which all vertices have even degree (cf. [26, Exercise 5.19]). Let us call such a subgraph an *even* subgraph, and an *odd* subgraph is defined similarly. Hence, every graph G contains an even induced subgraph with at least $|V(G)|/2$ vertices. The analogous properties for odd subgraphs seem to be more elusive. For a graph G, let $\mathsf{mos}(G)$ and $\chi_{\mathsf{odd}}(G)$ be the order of a largest odd induced subgraph of G and the minimum number of odd induced subgraphs of G that partition $V(G)$, respectively. Note that for $\chi_{\mathsf{odd}}(G)$ to be well-defined, each connected component of G must have even order.

Concerning the former parameter, the following long-standing –and still open– conjecture is cited as "part of the graph theory folklore" by Caro [7]: there exists a positive constant c such that every graph G without isolated vertices

* Work supported by French projects DEMOGRAPH (ANR-16-CE40-0028) and ESIGMA (ANR-17-CE23-0010), the program "Exploration Japon 2017" of the French embassy in Japan, and the JSPS KAKENHI grant number JP18K11157.

satisfies $\mathsf{mos}(G) \geq c \cdot |V(G)|$. In the following discussion we only consider graphs without isolated vertices. Caro [7] proved that $\mathsf{mos}(G) \geq (1 - o(1))\sqrt{n/6}$ where $n = |V(G)|$, and Scott [33] improved this bound to $\frac{cn}{\log n}$ for some $c > 0$. The conjecture has been proved for particular graph classes, such as trees [30], graphs of bounded chromatic number [33], graphs of maximum degree three [2], and graphs of tree-width at most two [20], also obtaining best possible constants.

As for the complexity of computing $\mathsf{mos}(G)$, Cai and Yang [6] studied, among other problems, two parameterized versions of this problem, and their reductions imply that it is NP-hard. They also prove the NP-hardness of computing the largest size of an even induced subgraph of a graph G, denoted $\mathsf{mes}(G)$. As a follow-up of [6], related problems were studied by Cygan et al. [9] and Goyal et al. [19].

The parameter χ_{odd}, which we call the *odd chromatic number*, has attracted much less interest in the literature. To the best of our knowledge, it has only been considered by Scott [34], who defined it (using a different notation) and proved that the necessary condition discussed above for $\chi_{\mathsf{odd}}(G)$ to be well-defined is also sufficient. He also provided lower and upper bounds on the maximum value of $\chi_{\mathsf{odd}}(G)$ over all n-vertex graphs. In particular, there are graphs G for which $\chi_{\mathsf{odd}}(G) = \Omega(\sqrt{n})$.

Our contribution. In this article we mostly focus on computational aspects of the parameters mos and χ_{odd}. Note that, given a graph G, deciding whether $\chi_{\mathsf{odd}}(G) \leq 1$ is trivial. We prove that deciding whether $\chi_{\mathsf{odd}}(G) \leq q$ is NP-complete for every $q \geq 3$ using a reduction from q-Coloring. We obtain a dichotomy on the complexity of computing χ_{odd} by showing that deciding whether $\chi_{\mathsf{odd}}(G) \leq 2$ can be solved in polynomial time, through a reduction to the existence of a feasible solution to a system of linear equations over GF[2].

Given the NP-hardness of computing both parameters, we are interested in its parameterized complexity [8, 11], namely in identifying relevant parameters k that allow for FPT algorithms, that is, algorithms running in time $f(k) \cdot n^{\mathcal{O}(1)}$ for some computable function f. Since the natural parameter, that is, the solution size, for mos has been studied by Cai and Yang [6] (and its dual as well), and for χ_{odd} the problem is para-NP-hard by our hardness results, we rather focus on *structural parameters*. Two of the most successful ones are definitely tree-width and clique-width, or its parametrically equivalent parameter *rank-width* introduced by Oum and Seymour [29]. This latter parameter is *stronger* than tree-width, in the sense that graph classes of bounded tree-width also have bounded rank-width. We present algorithms running in time $2^{\mathcal{O}(\mathsf{rw})} \cdot n^{\mathcal{O}(1)}$ for computing $\mathsf{mes}(G)$ and $\mathsf{mos}(G)$ for an n-vertex graph G given along with a decomposition tree of width at most rw, and an algorithm in time $2^{\mathcal{O}(q \cdot \mathsf{rw})} \cdot n^{\mathcal{O}(1)}$ for deciding whether $\chi_{\mathsf{odd}}(G) \leq q$. These algorithms are inspired by the ones of Bui-Xuan et al. [3, 4] to solve Maximum Independent Set parameterized by rank-width and boolean-width, respectively. To the best of our knowledge, our algorithms are the first ones parameterized by rank-width for an NP-hard problem running in time $2^{o(\mathsf{rw}^2)} \cdot n^{\mathcal{O}(1)}$ [1, 3, 17, 18, 28].

We also show that the dependency on rank-width of the above algorithms is asymptotically optimal under the Exponential Time Hypothesis (ETH) of Impagliazzo et al. [21, 22]. For this, it suffices to obtain a *linear* NP-hardness reduction from a problem for which a subexponential algorithm does not exist under the ETH. While our reduction to decide whether $\chi_{\mathsf{odd}}(G) \leq q$ already satisfies this property, the NP-hardness proof of Cai and Yang [6] for computing $\mathsf{mes}(G)$ and $\mathsf{mos}(G)$, which is from the Exact Odd Set problem [12], has a *quadratic* blow-up, so only a lower bound of $2^{o(\sqrt{n})}$ can be deduced from it. Motivated by this, we present linear NP-hardness reductions from 2in3-Sat to the problems of computing $\mathsf{mes}(G)$ and $\mathsf{mos}(G)$. The reduction itself is not very complicated, but the correctness proof requires some non-trivial arguments[3].

Finally, motivated by the complexity of computing these parameters, we obtain two tight bounds on their values. We first prove that for every graph G with all components of even order, $\chi_{\mathsf{odd}}(G) \leq \mathsf{tw}(G)+1$, where $\mathsf{tw}(G)$ denotes the tree-width of G. This result improves the best known lower bound on a parameter defined by Hou et al. [20] (cf. Section 5 for the details). On the other hand, we prove that, for every n-vertex graph G such that $V(G)$ can be partitioned into two non-empty sets that are complete to each other (i.e., a *join*), $\mathsf{mos}(G) \geq 2 \cdot \lceil \frac{n-2}{4} \rceil$. In particular, this proves the conjecture about the linear size of an odd induced subgraph for *cographs*, which are the graphs of clique-width two. This adds another graph class to the previous ones for which the conjecture is known to be true [2, 20, 30, 33]. It is interesting to mention that our proof implies that, for a cograph G, $\chi_{\mathsf{odd}}(G) \leq 3$, and this bound is also tight. While for cographs, or equivalently P_4-free graphs, we have proved that the odd chromatic number is bounded, we also show that it is unbounded for P_5-free graphs.

Organization. We start with some preliminaries in Section 2. In Section 3 we provide the linear NP-hardness reductions and the polynomial-time algorithm for deciding whether $\chi_{\mathsf{odd}}(G) \leq 2$. The FPT algorithms by rank-width are presented in Section 4, and the tight bounds in Section 5. We conclude the article in Section 6 with a number of open problems and research directions. Additional results for related problems can be found in the full version, available at `https://arxiv.org/abs/2002.06078`. Due to space limitations, the proofs of the results marked with '($\star$)' can be found in the full version.

2 Preliminaries

Graphs. We use standard graph-theoretic notation, and we refer the reader to [10] for any undefined notation. Let $G = (V, E)$ be a graph, $S \subseteq V$, and H be a subgraph of G. We denote an edge between u and v by uv. The *order* of G is $|V|$. The *degree* (resp. *open neighborhood*, *closed neighborhood*) of a vertex $v \in V$ is denoted by $\mathsf{deg}(v)$ (resp. $N(v)$, $N[v]$), and we let $\mathsf{deg}_H(v) = |N(v) \cap V(H)|$.

[3] We would like to mention that another NP-hardness proof for computing $\mathsf{mes}(G)$ has very recently appeared online [32]. The proof uses a chain of reductions from Maximum Cut and, although it also involves a quadratic blow-up, it can be avoided by starting from Maximum Cut restricted to graphs of bounded degree.

We use the notation $G - S = G[V(G) \setminus S]$. The *maximum* and *minimum degree* of G are denoted by $\Delta(G)$ and $\delta(G)$, respectively. We denote by P_i the path on i vertices. For two graphs G_1 and G_2, with $V(G_2) \subseteq V(G_1)$, the *union* of G_1 and G_2 is the graph $(V(G_1), E(G_1) \cup E(G_2))$. The operation of *contracting* an edge uv consists in deleting both u and v and adding a new vertex w with neighbors $N(u) \cup N(v) \setminus \{u, v\}$. A graph M is a *minor* of G if it can be obtained from a subgraph of G by a sequence of edge contractions. For a positive integer $k \geq 3$, the *k-wheel* is the graph obtained from a cycle C on k vertices by adding a new vertex v adjacent to all the vertices of C. A *join* in a graph G is a partition of $V(G)$ into two non-empty sets V_1 and V_2 such that every vertex in V_1 is adjacent to every vertex in V_2. For a positive integer i, we denote by $[i]$ the set containing every integer j such that $1 \leq j \leq i$.

Parameterized complexity. We refer the reader to [8, 11, 14, 27] for basic background on parameterized complexity, and we recall here only some basic definitions. A *parameterized problem* is a decision problem whose instances are pairs $(x, k) \in \Sigma^* \times \mathbb{N}$, where k is called the *parameter*. A parameterized problem is *fixed-parameter tractable* (FPT) if there exists an algorithm $\mathcal{A}$, a computable function f, and a constant c such that given an instance $I = (x, k)$, $\mathcal{A}$ (called an FPT *algorithm*) correctly decides whether $I \in L$ in time $f(k) \cdot |I|^c$. A parameterized problem is *slice-wise polynomial* (XP) if there exists an algorithm $\mathcal{A}$ and two computable functions f, g such that given an instance $I = (x, k)$, $\mathcal{A}$ (called an XP *algorithm*) correctly decides whether $I \in L$ in time $f(k) \cdot |I|^{g(k)}$.

The *Exponential Time Hypothesis* (ETH) of Impagliazzo et al. [21, 22] implies that the 3-Sat problem on n variables cannot be solved in time $2^{o(n)}$. We say that a polynomial reduction from a problem Π_1 to a problem Π_2, generating an input of size n_2 from an input of size n_1, is *linear* if $n_2 = \mathcal{O}(n_1)$. Clearly, if Π_1 cannot be solved, under the ETH, in time $2^{o(n)}$ on inputs of size n, and there exists a linear reduction from Π_1 to Π_2, then Π_2 cannot either.

Width parameters. In this article we mention several width parameters of graphs, such as tree-width, rank-width, clique-width, or boolean-width. However, since we only deal with rank-width in our algorithms (cf. Section 4), we give only the definition of this parameter here.

A *decomposition tree* of a graph G is a pair (T, δ) where T is a full binary tree (i.e., T is rooted and every non-leaf node has two children) and δ a bijection between the leaf set of T and the vertex set of G. For a node w of T, we denote by V_w the subset of $V(G)$ in bijection –via δ– with the leaves of the subtree of T rooted at w. We say that the decomposition defines the *cut* $\left(V_w, \overline{V_w}\right)$. The *rank-width* of a decomposition tree (T, δ) of a graph G, denoted by $\mathsf{rw}(T, \delta)$, is the maximum over all $w \in V(T)$ of the rank of the adjacency matrix of the bipartite graph $G[V_w, \overline{V_w}]$. The *rank-width of* G, denoted by $\mathsf{rw}(G)$, is the minimum $\mathsf{rw}(T, \delta)$ over all decomposition trees (T, δ) of G.

Definition of the problems. A graph is called *odd* (resp. *even*) if every vertex has odd (resp. even) degree. The Maximum Odd Subgraph (resp. Maximum Even Subgraph problem consists in, given a graph G, determining the maximum order of an odd (resp. even) induced subgraph of G, that is, $\mathsf{mos}(G)$ (resp. $\mathsf{mes}(G)$). An *odd q-coloring* of a graph $G = (V, E)$ is a set of q odd induced subgraphs $H_1, \dots, H_q$ of G such that $V(H_1) \uplus \cdots \uplus V(H_q)$ is a partition of V. The Odd q-Coloring problem consists in determining whether an input graph G admits an odd q-coloring. In the Odd Chromatic Number problem, the objective is to determine the smallest integer q such that an input graph G admits an odd q-coloring.

3 Linear reductions and a polynomial-time algorithm

We first present the linear reductions for Maximum Even Subgraph and Maximum Odd Subgraph, and then for Odd q-Coloring for $q \geq 3$.

Theorem 1 ($\star$). *The* Maximum Even Subgraph *and* Maximum Odd Subgraph *problems are* NP*-hard. Moreover, none of them can be solved in time $2^{o(n)}$ on n-vertex graphs unless the* ETH *fails.*

Theorem 2. *For every integer $q \geq 3$, given a graph G on n vertices, determining whether $\chi_{\mathsf{odd}}(G) \leq q$ is* NP*-complete and, moreover, cannot be solved in time $2^{o(n)}$ unless the* ETH *fails.*

Proof: Membership in NP is clear. For every integer $q \geq 3$, we present a linear reduction from the q-Coloring problem, which is well-known to be NP-hard and not solvable in time $2^{o(n)}$ on n-vertex graphs unless the ETH fails [21,22]. We will use the fact that any graph $G = (V, E)$ such that $|V| + |E|$ is even admits an orientation of E such that, in the resulting digraph, all the vertex in-degrees are odd; we call such an orientation an *odd orientation*. Moreover, an odd orientation can be found in polynomial time (for a proof, see for instance [16]).

Given an instance $G = (V, E)$ of q-Coloring, we build from G an instance $G^\bullet$ of Odd q-Coloring as follows. First, if $|V| + |E|$ is odd, we arbitrarily select a vertex $v \in V$ and add a triangle on three new vertices v_1, v_2, v_3 and the edge vv_1. Note that the resulting graph $G' = (V', E')$ is q-colorable for $q \geq 3$ if and only if G is, and that $|V'| + |E'|$ is even. Hence, E' admits an odd orientation ϕ. We let $G^\bullet$ be the graph obtained from G' by subdividing every edge once. Note that the size of $G^\bullet$ depends linearly on the size of G, as required. We claim that $\chi(G) \leq q$ if and only if $\chi_{\mathsf{odd}}(G^\bullet) \leq q$.

Assume first that we are given a proper q-coloring $c : V \to [q]$, which can trivially be extended to a proper q-coloring of G'. We define an odd q-coloring c_{odd} of $G^\bullet$ as follows. If $v \in V(G^\bullet)$ is an original vertex of V', we set $c_{\mathsf{odd}}(v) = c(v)$. Otherwise, if v is a subdivision vertex between two vertices u and w of V', we set $c_{\mathsf{odd}}(v) = c(u)$ if edge uw is oriented toward u in ϕ, and $c_{\mathsf{odd}}(v) = c(w)$ otherwise. It can be easily verified that c_{odd} is indeed an odd q-coloring of $G^\bullet$.

Conversely, let $c_{\mathsf{odd}} : V(G^\bullet) \to [q]$ be an odd q-coloring of $G^\bullet$, let uw be an edge of G', and let v be the subdivision vertex in $G^\bullet$ between u and w. If follows that $c_{\mathsf{odd}}(u) \neq c_{\mathsf{odd}}(w)$, as otherwise vertex v would have degree zero or two in its color class. Therefore, letting $c(v) = c_{\mathsf{odd}}(v)$ for every vertex $v \in V(G)$ defines a proper q-coloring of G, and the theorem follows. $\square$

Theorem 2 establishes the NP-hardness of Odd q-Coloring for every $q \geq 3$. On the other hand, the Odd 1-Coloring is trivial, as for any graph G, $\chi_{\mathsf{odd}}(G) \leq 1$ if and only if G is an odd graph itself. Therefore, the only remaining case is Odd 2-Coloring. In the next theorem we prove that this problem can be solved in polynomial time.

Theorem 3. *The* Odd 2-Coloring *problem can be solved in polynomial time.*

Proof: We will express the Odd 2-Coloring problem as the existence of a feasible solution to a system of linear equations over the binary field, which can be determined in polynomial time using, for instance, Gaussian elimination. Given an instance $G = (V, E)$ of Odd 2-Coloring, let its vertices be labeled $v_1, \ldots, v_n$. For every vertex $v_i \in V$ we create a binary variable x_i, and for every edge $v_i v_j \in E$, we create a binary variable $x_{i,j}$. The interpretation of these two types of variables is quite different. Namely, for a vertex variable x_i, its value corresponds to the color (either 0 or 1) assigned to vertex v_i. On the other hand, the value of an edge variable corresponds the whether this edge belongs to a monochromatic subgraph, that is, to whether both its endvertices get the same color. In this case, its value is 1, and 0 otherwise. We guarantee this latter property by adding the following set of linear equations:

$$x_i + x_j + x_{i,j} \equiv 1 \qquad \text{for every edge } v_i v_j \in E. \tag{1}$$

To guarantee that the degree of every vertex in each of the two monochromatic subgraphs is odd, we add the following set of linear equations (for an edge variable $x_{i,j}$, to simplify the notation we interpret $x_{j,i} = x_{i,j}$):

$$\sum_{j : v_j \in N(v_i)} x_{i,j} \equiv 1 \qquad \text{for every vertex } v_i \in V. \tag{2}$$

Note that by Equation (1), only monochromatic edges contribute to the sum of Equation (2). Therefore, the above discussion implies that $\chi_{\mathsf{odd}}(G) \leq 2$ if and only if the system of linear equations given by Equations (1) and (2) admits a feasible solution, and the theorem follows. □

Note that the Even 2-Coloring problem could be formulated in a similar way, just by replacing Equation (2) with $\sum_{j : v_j \in N(v_i)} x_{i,j} \equiv 0$. However, this is not that interesting, since all the instances of Even 2-Coloring are positive [26].

4 Dynamic programming algorithms

In this section, we present FPT algorithms for Maximum Odd/Even Subgraph and Odd q-Coloring, parameterized by the rank-width of the input graph. The algorithms are similar to those of Bui-Xuan et al. [3, 4] for Maximum Independent Set parameterized by rank-width and boolean-width, respectively, and also to the one by Bui-Xuan et al. [5] for so-called locally checkable vertex partitioning problems. There are however two key differences with our algorithms.

First, while partial solutions for Maximum Independent Set are, themselves, independent sets, this is not true in general for odd subgraphs, where partial solutions may consist in a subgraph some vertices of which have even degree. Those vertices will impose some extra constraints on the remainder of the solution. The second difference is that, while the equivalence classes of [3] and [4] are based on neighborhoods of vertex sets, those for Maximum Odd Subgraph only require "neighborhoods modulo 2". This will allow us to consider only $2^{\mathcal{O}(\mathsf{rw})}$ equivalence classes, compared to $2^{\mathcal{O}(\mathsf{rw}^2)}$ classes used in [3] for Maximum Independent Set.

Throughout this section, we will rely on the notion of "neighborhood modulo 2" of a set of vertices, defined as follows. Given a graph G and $X \subseteq V(G)$, the *neighborhood of X modulo 2*, denoted by $N_2(X)$, is the set $\triangle_{u\in X}(N(u))$, where the operator $\triangle$ denotes the symmetric difference. Note that $N_2(X)$ is exactly the set of vertices in $V(G) \setminus X$ that have an odd number of neighbors in X. The results in this section are stated using the $\mathcal{O}^*$ notation, which hides polynomial factors in the input size.

Theorem 4. *Given a graph G along with a decomposition tree of rank-width* rw*, the* Maximum Odd Subgraph *problem can be solved in time $\mathcal{O}^*(2^{3\mathsf{rw}})$.*

Proof: We give a dynamic programming over the given decomposition tree (T, L). Recall that there is a bijection between the leaves of T and $V(G)$, and that each edge of T corresponds to a cut $(A, \overline{A})$ of G. We begin by defining the equivalence relation over subsets of A, given a cut $(A, \overline{A})$: two sets $X, Y \subseteq V(G)$ are *odd neighborhood equivalent* with regard to A, denoted by $X \equiv_2^A Y$, if $N_2(X) \setminus A = N_2(Y) \setminus A$. Then, given a row basis $\mathcal{B}$ of the adjacency matrix of $(A, \overline{A})$ over $\mathsf{GF}[2]$, where we interpret a vertex set as the vector corresponding to its vertices, we define the *representative* of a set $X \subseteq A$ as the the unique set of vertices $R_A(X) \subseteq A$ such that $R_A(X) \subseteq \mathcal{B}$ and $X \equiv_2^A R_A(X)$. Observe that since $(A, \overline{A})$ is a cut of (T, L), its adjacency matrix has rank at most $\mathsf{rw}(G)$, and therefore $|R_A(X)| \leq \mathsf{rw}(G)$. This implies, in particular, that there are at most $2^{\mathsf{rw}(G)}$ distinct representatives for subsets of a given set A.

We are now ready to define the tables of our algorithm. Given an edge e of (T, L) and its associated cut $(A, \overline{A})$ of G, we store in table T_A, for every pair R, R' of representatives of subsets of A and $\overline{A}$, respectively, a largest set $S \subseteq A$ such that S is odd neighborhood equivalent to R, and all the vertices that have even degree in $G[S]$ is exactly the set $N_2(R') \cap S$. More formally:

$$(\maltese)\ T_A[R, R'] = \operatorname*{maxset}_{S\subseteq A}\{S \equiv_2^A R \wedge \{v \in S : |N(v) \cap S| \text{ is even}\} = N_2(R') \cap S\},$$

where the notation 'maxset' indicates a largest set that satisfies the conditions. In cases where edge e is incident with a leaf, the cut associated with e is of the form $(\{u\}, V(G)\setminus\{u\})$. We set $T_{\{u\}}[\emptyset, \emptyset] = T_{\{u\}}[\emptyset, \{v\}] = \emptyset$, and $T_{\{u\}}[\{u\}, \{v\}] = \{u\}$, where v is the unique vertex of a basis of the adjacency matrix of the cut $(V(G) \setminus \{u\}, \{u\})$, which is the only non-empty choice for R'. The entry $T_{\{u\}}[\{u\}, \emptyset]$ is left empty, due to there being no subgraph of $G[\{u\}]$ with the same neighborhood as $\{u\}$ in $G - \{u\}$, all vertices of which that have even degree lying in $N_2(\emptyset)$.

Given an edge e of (T, L) such that the tables of both edges incident with one endvertex of e, say f, f', have been computed, we compute the table of e as follows. Let us denote by $(A, \overline{A}), (X, \overline{X})$, and $(Y, \overline{Y})$ the cuts associated with e, f, and f', respectively. For each pair of representatives $R_A, R_{\overline{A}}$ of the cut $(A, \overline{A})$, the value of $T_A[R_A, R_{\overline{A}}]$ is the largest $T_X[R_X, R_{\overline{X}}] \cup T_Y[R_Y, R_{\overline{Y}}]$, such that $R_X, R_{\overline{X}}, R_Y$, and $R_{\overline{Y}}$ satisfy the following conditions with regard to R_A and $R_{\overline{A}}$:

(i) $R_A \equiv_2^A R_X \triangle R_Y$, (ii) $R_{\overline{X}} \equiv_2^{\overline{X}} R_{\overline{A}} \triangle R_Y$, and (ii') $R_{\overline{Y}} \equiv_2^{\overline{Y}} R_{\overline{A}} \triangle R_X$.

We proceed with this computation, starting from the leaves, in a bottom-up manner, having previously rooted T by choosing an arbitrary edge, subdividing it, and making the newly created vertex the root of T. Observe that in the final stage of the algorithm, when the tables of both edges f, f' incident with the root have been computed, we compute the table for the root node as described above, with $\overline{A} = \emptyset$, since $X \cup Y = V(G)$ in this case. Of the three conditions described above, condition (i) becomes trivial, since $R_A = \emptyset$, and conditions (ii) and (ii') simplify to $R_{\overline{X}} \equiv_2^{\overline{X}} R_Y$, and $R_{\overline{Y}} \equiv_2^{\overline{Y}} R_X$, respectively.

We first observe that since, as noted above, there are at most 2^{rw} representatives on each side of each cut, and the choices of R_X, R_Y and $R_{\overline{A}}$ uniquely determines $R_{\overline{X}}, R_{\overline{Y}}$ and R_A through equations (i), (ii), and (ii'), and computing new tables can be carried out in time $\mathcal{O}^*(2^{3\mathsf{rw}})$, as desired. It now remains to prove that the algorithm correctly computes an optimal solution. The correctness of the tables for the leaves of T follows from their description. We now prove by induction that the tables are correct for internal edges of T as well. Let us assume T_X and T_Y have been fully and correctly computed for all possible representatives $R_X, R_{\overline{X}}, R_Y$, and $R_{\overline{Y}}$ as per the description above. We first argue that the tables' description is correct, i.e., given an optimal solution OPT (that is, an induced subgraph of G achieving $\mathsf{mos}(G)$) and a cut $(A, \overline{A})$, $S = \mathsf{OPT} \cap A$ is a largest set that satisfies (✠) for some pair R, R' of representatives. Indeed, assume for contradiction that there exists $S^* \subseteq A$ such that $S^* \equiv_2^A S$, $\{v \in S^* : |N(v) \cap S^*| \text{ is even}\} = S^* \cap N_2(\mathsf{OPT} \cap \overline{A})$, and $|S| < |S^*|$. Then, $\mathsf{OPT}^* = (\mathsf{OPT} \setminus S) \cup S^*$ induces an odd subgraph of G and $|\mathsf{OPT}^*| > |\mathsf{OPT}|$, contradicting the optimality of OPT.

Finally, we argue that if T_X and T_Y are computed correctly, then so is T_A, i.e., given any two representatives R_A and $R_{\overline{A}}$ of A and $\overline{A}$, respectively, there exist representatives $R_X, R_{\overline{X}}, R_Y$, and $R_{\overline{Y}}$ of $X, \overline{X}, Y$, and $\overline{Y}$, respectively, that satisfy conditions (i), (ii), and (ii'), and such that $T_X[R_X, R_{\overline{X}}] \cup T_Y[R_Y, R_{\overline{Y}}]$ is a largest set that satisfies (✠) with respect to $(R_A, R_{\overline{A}})$. Let $R_X, R_{\overline{X}}, R_Y$, and $R_{\overline{Y}}$ be representatives such that $T_A[R_A, R_{\overline{A}}] = T_X[R_X, R_{\overline{X}}] \cup T_Y[R_Y, R_{\overline{Y}}] = S$, and let S_X and S_Y denote $S \cap X$ and $S \cap Y$, respectively. Note that, since X and Y form a partition of A, S_X and S_Y form a partition of S, which implies $S = S_X \cup S_Y = S_X \triangle S_Y$. We first show that S indeed satisfies (✠) with respect to $(A, \overline{A})$, i.e., $S \equiv_2^A R_A$ and $\{v \in S : |N(v) \cap S| \text{ is even}\} = N_2(R_{\overline{A}}) \cap S$. For the first of those two conditions, combining it with the fact that $S = S_X \cup S_Y = S_X \triangle S_Y$, we only need to prove that $S_X \triangle S_Y \equiv_2^A R_X \triangle R_Y$. Observe first that, since X and Y form a partition of A, we have that for every vertex $v \in \overline{A}$, $|N(v) \cap A| = |N(v) \cap X| + |N(v) \cap Y|$. Therefore, for every sets $X', X'' \subseteq X$ and $Y', Y'' \subseteq Y$, it holds that if $X' \equiv_2^X X''$ and $Y' \equiv_2^Y Y''$, then $X' \triangle Y' \equiv_2^A X'' \triangle Y''$. From the definition of representative we obtain that $S \equiv_2^A S_X \triangle S_Y \equiv_2^A R_X \triangle R_Y$, as desired.

Let us now consider the second condition, i.e., $\{v \in S : |N(v) \cap S| \text{ is even}\} = N_2(R_{\overline{A}}) \cap S$. Let us assume first that $v \in S_X$. If $|N(v) \cap S|$ is even, then at least one of the following cases holds:

- $|N(v) \cap S_X|$ is even and $v \notin N_2(S_Y)$. Since $|N(v) \cap S_X|$ is even, we obtain from (✠) in T_X that $v \in N_2(R_{\overline{X}})$, which when combined with (ii) implies $v \in N_2(\overline{A}) \triangle N_2(S_Y)$. Since $v \notin N_2(S_Y)$, it follows that $v \in N_2(\overline{A})$, as desired.
- $|N(v) \cap S_X|$ is odd and $v \in N_2(S_Y)$. Symmetrically to the case above, we have that $v \notin N_2(R_{\overline{X}})$, hence $v \notin N_2(\overline{A}) \triangle N_2(S_Y)$ from (ii), and since $v \in N_2(S_Y)$, it follows that $v \in N_2(\overline{A})$, as desired.

The case where $v \in S_Y$ is proved similarly, replacing condition (ii) with (ii'). Therefore, $\{v \in S : |N(v) \cap S| \text{ is even}\} \subseteq S \cap N_2(R_{\overline{A}})$. Let us now assume that $v \in S_X \cap N_2(R_{\overline{A}})$. From (ii), we obtain that $v \in N_2(S \cap \overline{X})$ if and only if $v \notin N_2(S_Y)$. Since T_X satisfies (✠), it holds that $v \in N_2(S \cap \overline{X})$ if and only if $|N(v) \cap S_X|$ is even, and therefore $v \notin N_2(S_Y)$ if and only if $|N(v) \cap S_X|$ is even. Therefore, $|N(v) \cap S| = |N(v) \cap S_X| + |N(v) \cap S_Y|$ is even, as desired. As above, the case where $v \in S_Y$ is proved similarly, replacing condition (ii) with (ii'). Therefore, $\{v \in S : |N(v) \cap S| \text{ is even}\} = S \cap N_2(R_{\overline{A}})$.

Finally, we prove the maximality of S among all those sets that satisfy (✠) with respect to $(R_A, R_{\overline{A}})$. Let us assume for a contradiction that there exists S^* that satisfies (✠) with respect to $(R_A, R_{\overline{A}})$ and such that $|S^*| > |S|$. Let S^*_X and S^*_Y denote $S^* \cap X$ and $S^* \cap Y$, respectively. Observe that S^*_X and S^*_Y satisfy (✠) with respect to some pairs of representatives $(R_X, R_{\overline{X}})$ and $(R_Y, R_{\overline{Y}})$, respectively. In addition, observe that since S satisfies (✠) with respect to $(R_A, R_{\overline{A}})$, it follows that S, S_X, and S_Y satisfy conditions (i), (ii), and (ii') with respect to $(R_A, R_{\overline{A}})$, contradicting the assumption that T_X and T_Y were computed correctly. □

Small variations of Theorem 4 allow us to prove the following two theorems.

Theorem 5 ($\star$). *Given a graph G along with a decomposition tree of rank-width* rw*, the* Maximum Even Subgraph *problem can be solved in time $\mathcal{O}^*(2^{3\mathsf{rw}})$.*

Theorem 6 ($\star$). *Given a graph G along with a decomposition tree of rank-width w, the* Odd q-Coloring *problem can be solved in time $\mathcal{O}^*(2^{\mathcal{O}(q \cdot \mathsf{rw})})$.*

5 Tight bounds

In this section we provide two tight bounds concerning odd induced subgraphs and odd colorings. Namely, we first provide in Theorem 7 a tight upper bound on the odd chromatic number in terms of tree-width, and then we provide in Theorem 8 a tight lower bound on the size of a maximum odd induced subgraph for graphs that admit a join.

Theorem 7. *For every graph G with all components of even order we have that $\chi_{\mathsf{odd}}(G) \leq \mathsf{tw}(G)+1$, and this bound is tight.*

Proof: Scott proved [34, Corollary 3] that every graph G with all components of even order admits a vertex partition such that every vertex class induces a tree with all degrees odd. Consider such a vertex partition, and let G' be the graph obtained from G by contracting each of the trees to a single vertex. Since G' is a minor of G, we have that $\mathsf{tw}(G') \leq \mathsf{tw}(G)$. Now note that every proper vertex coloring of G' using q colors can be lifted to a partition of $V(G)$ into q odd induced subgraphs (in fact, odd induced forests). Indeed, with every color i of a proper q-coloring of $V(G')$ we associate an induced forest of G defined by the union of the trees whose corresponding vertex in G' is colored i. Therefore,

$$\chi_{\mathsf{odd}}(G) \leq \chi(G') \leq \mathsf{tw}(G')+1 \leq \mathsf{tw}(G)+1,$$

where we have used the well-known fact that the chromatic number of a graph is at most its tree-width plus one [23].

To see that this bound it tight, consider a subdivided clique $K_n^{\bullet}$, that is, the graph obtained from a clique on n vertices, with $n \equiv 0,3 \pmod 4$, by subdividing every edge once. Since no pair of original vertices of the clique can get the same color, we have that $\chi_{\mathsf{odd}}(K_n^{\bullet}) = n = \mathsf{tw}(K_n^{\bullet})+1$. □

Let us mention some consequences of Theorem 7. Hou et al. [20] define the following parameter. Let $\mathcal{G}_k$ be the set of all graphs of treewidth at most k without isolated vertices, and let $c_k = \min_{G \in \mathcal{G}_k} \frac{\mathsf{mos}(G)}{|V(G)|}$. In [20] the authors prove that $c_2 = 2/5$ and say that the best general lower bound is $c_k \geq \frac{1}{2k+2}$, which follows from a result of Scott [33]. As an immediate corollary of Theorem 7 it follows that $c_k \geq \frac{1}{k+1}$, which improves the lower bound by a factor two. As it is known [20] that, for $k \in [4]$, $c_k \leq \frac{2}{k+3}$, our lower bound implies that $1/4 \leq c_3 \leq 1/3$ and $1/5 \leq c_4 \leq 2/7$.

We now provide a lower bound on $\mathsf{mos}(G)$ for every graph that admits a join.

Theorem 8 ($\star$). *For every n-vertex graph G that admits a join we have*

$$\mathsf{mos}(G) \geq 2 \cdot \left\lceil \frac{n-2}{4} \right\rceil, \text{ and this bound is tight even for cographs.}$$

Determining a tight lower bound for cographs that are not necessarily connected remains open. The proof of Cases 1 and 2 of Theorem 7 together with the fact that $\chi_{\mathsf{odd}}(K_{2,2,2}) = 3$ (since $\mathsf{mos}(K_{2,2,2}) = 2$) yield the following corollary.

Corollary 1. *Let G be a cograph with every connected component of even order. Then $\chi_{\mathsf{odd}}(G) \leq 3$. Moreover, this bound is tight.*

Note that cographs can be equivalently defined as P_4-free graphs. It is interesting to note that, in contrast to Corollary 1, P_5-free graphs have unbounded odd chromatic number. Indeed, let H_n be the graph obtained from the subdivided clique $K_n^\bullet$, with $n \equiv 0, 3 \pmod 4$, by adding an edge between each pair of original vertices of the clique. It can be checked that $\chi_{\mathsf{odd}}(H_n) \geq n$ and, in fact, the proof of Theorem 2 implies that $\chi_{\mathsf{odd}}(H_n) = n$. Note that H_n is a split graph, hence split graphs have unbounded odd chromatic number.

6 Further research

We considered computational aspects of the Maximum Odd Subgraph and Odd q-Coloring problems. A number of interesting questions remain open.

We gave in Theorem 6 an algorithm that solves Odd q-Coloring in time $\mathcal{O}^*(2^{\mathcal{O}(q \cdot \mathsf{rw})})$. Is the Odd Chromatic Number problem FPT or W[1]-hard parameterized by rank-width? A strongly related question is how the odd chromatic number depends on rank-width. We proved in Theorem 7 that $\chi_{\mathsf{odd}}(G) \leq \mathsf{tw}(G) + 1$, but we do not know whether $\chi_{\mathsf{odd}}(G) \leq f(\mathsf{rw}(G))$ for some function f. Note that this would not only yield an FPT algorithm for Odd Chromatic Number by rank-width, but would also prove the conjecture about the linear size of a largest odd induced subgraph [7] for all graphs of bounded rank-width. As a first step in this direction, we proved in Corollary 1 that cographs, which have rank-width at most one, have odd chromatic number at most three. It would be interesting to prove an upper bound for distance-hereditary graphs, which can be equivalently defined as graphs of rank-width one.

In fact, we do not even know whether Odd Chromatic Number by rank-width is in XP. In view of the algorithm of Theorem 6, a sufficient condition for this would be that there exists a function f such that $\chi_{\mathsf{odd}}(G) \leq f(\mathsf{rw}(G)) \cdot \log |V(G)|$ for every graph G with all components of even order. Another promising strategy would be to generalize the XP algorithms of Rao [31] to *counting* monadic second-odder logic. Toward an eventual W[1]-hardness proof, a natural strategy is to try to adapt the reduction given by Fomin et al. [15] to prove that Chromatic Number is W[1]-hard by clique-width (hence, rank-width). This reduction is from Equitable Coloring parameterized by the number of colors plus tree-width, proved to be W[1]-hard by Fellows et al. [13]. By appropriately modifying the chain of reductions given in [13], we have only managed to prove that the naturally defined Odd Equitable Coloring problem is W[1]-hard by tree-width, but not if we add the number of colors as a parameter.

Concerning Odd q-Coloring parameterized by tree-width, a straightforward dynamic programming algorithm that guesses, for every vertex, its color class and the parity of its degree within that class, runs in time $\mathcal{O}^*((2q)^{\mathsf{tw}})$. Note that this algorithm together with Theorem 7 yield an algorithm for Odd Chromatic Number in time $\mathcal{O}^*((2\mathsf{tw} + 2)^{\mathsf{tw}})$. By the lower bound under the ETH of Lokshtanov et al. [25] for Chromatic Number by tree-width and the fact that our reduction of Theorem 2 preserves tree-width, it follows that the dependency on tree-width of this algorithm is asymptotically optimal under the ETH. It would be interesting to prove lower bounds under the *Strong Exponential Time Hypothesis* (SETH). Note that our reduction of Theorem 2 together with the lower bound under the SETH of Lokshtanov et al. [24] for q-Coloring by tree-width yield a lower bound for Odd q-Coloring of $\mathcal{O}^*((q - \varepsilon)^{\mathsf{tw}})$ under the SETH.

References

1. B. Bergougnoux and M. M. Kanté. Rank based approach on graphs with structured neighborhood. *CoRR*, abs/1805.11275, 2018.
2. D. M. Berman, H. Wang, and L. Wargo. Odd induced subgraphs in graphs of maximum degree three. *Australasian Journal of Combinatorics*, 15:81–86, 1997.
3. B. Bui-Xuan, J. A. Telle, and M. Vatshelle. H-join decomposable graphs and algorithms with runtime single exponential in rankwidth. *Discrete Applied Mathematics*, 158(7):809–819, 2010.
4. B. Bui-Xuan, J. A. Telle, and M. Vatshelle. Boolean-width of graphs. *Theoretical Computer Science*, 412(39):5187–5204, 2011.
5. B. Bui-Xuan, J. A. Telle, and M. Vatshelle. Fast dynamic programming for locally checkable vertex subset and vertex partitioning problems. *Theoretical Computer Science*, 511:66–76, 2013.
6. L. Cai and B. Yang. Parameterized complexity of even/odd subgraph problems. *Journal of Discrete Algorithms*, 9(3):231–240, 2011.
7. Y. Caro. On induced subgraphs with odd degrees. *Discrete Mathematics*, 132(1-3):23–28, 1994.
8. M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.
9. M. Cygan, D. Marx, M. Pilipczuk, M. Pilipczuk, and I. Schlotter. Parameterized Complexity of Eulerian Deletion Problems. *Algorithmica*, 68(1):41–61, 2014.
10. R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.
11. R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
12. R. G. Downey, M. R. Fellows, A. Vardy, and G. Whittle. The parametrized complexity of some fundamental problems in coding theory. *SIAM Journal on Computing*, 29(2):545–570, 1999.
13. M. R. Fellows, F. V. Fomin, D. Lokshtanov, F. A. Rosamond, S. Saurabh, S. Szeider, and C. Thomassen. On the complexity of some colorful problems parameterized by treewidth. *Information and Computation*, 209(2):143–153, 2011.
14. J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. Springer, 2006.
15. F. V. Fomin, P. A. Golovach, D. Lokshtanov, and S. Saurabh. Intractability of Clique-Width Parameterizations. *SIAM Journal on Computing*, 39(5):1941–1956, 2010.
16. A. Frank, T. Jordán, and Z. Szigeti. An orientation theorem with parity conditions. In *Proc. of the 7th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 1610 of *LNCS*, pages 183–190, 1999.
17. R. Ganian, P. Hlinený, and J. Obdrzálek. Better algorithms for satisfiability problems for formulas of bounded rank-width. *Fundamenta Informaticae*, 123(1):59–76, 2013.
18. R. Ganian, P. Hlinený, and J. Obdrzálek. A unified approach to polynomial algorithms on graphs of bounded (bi-)rank-width. *European Journal of Combinatorics*, 34(3):680–701, 2013.
19. P. Goyal, P. Misra, F. Panolan, G. Philip, and S. Saurabh. Finding even subgraphs even faster. *Journal of Computer and System Sciences*, 97:1–13, 2018.
20. X. Hou, L. Yu, J. Li, and B. Liu. Odd induced subgraphs in graphs with treewidth at most two. *Graphs and Combinatorics*, 34(4):535–544, 2018.

21. R. Impagliazzo and R. Paturi. On the Complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
22. R. Impagliazzo, R. Paturi, and F. Zane. Which Problems Have Strongly Exponential Complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001.
23. T. Kloks. *Treewidth. Computations and Approximations*. Springer-Verlag LNCS, 1994.
24. D. Lokshtanov, D. Marx, and S. Saurabh. Known algorithms on graphs of bounded treewidth are probably optimal. *ACM Transactions on Algorithms*, 14(2):13:1–13:30, 2018.
25. D. Lokshtanov, D. Marx, and S. Saurabh. Slightly superexponential parameterized problems. *SIAM Journal on Computing*, 47(3):675–702, 2018.
26. L. Lovász. *Combinatorial Problems and Exercises*. North-Holland, Amsterdam, 1979.
27. R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.
28. S. Oum. Rank-width: Algorithmic and structural results. *Discrete Applied Mathematics*, 231:15–24, 2017.
29. S. Oum and P. D. Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96(4):514–528, 2006.
30. A. J. Radcliffe and A. D. Scott. Every tree contains a large induced subgraph with all degrees odd. *Discrete Mathematics*, 140(1-3):275–279, 1995.
31. M. Rao. MSOL partitioning problems on graphs of bounded treewidth and clique-width. *Theoretical Computer Science*, 377(1-3):260–267, 2007.
32. A. Röyskö. https://cstheory.stackexchange.com/questions/45885/complexity-of-finding-the-largest-induced-subgraph-with-all-even-degrees.
33. A. D. Scott. Large induced subgraphs with all degrees odd. *Combinatorics, Probability & Computing*, 1:335–349, 1992.
34. A. D. Scott. On induced subgraphs with all degrees odd. *Graphs and Combinatorics*, 17(3):539–553, 2001.