

Original citation:

Gajarsky, Jakub, Hlineny, Petr , Lokshtanov, Daniel , Obdrálek, Jan, Ordyniak, Sebastian, Ramanujan, Maadapuzhi Sridharan and Saurabh, Saket (2015) FO model checking on posets of bounded width. In: IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015, Berkeley, CA, USA, 17-20 Oct 2015. Published in: 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS) pp. 963-974.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/97713>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting /republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

FO Model Checking on Posets of Bounded Width

Jakub Gajarský³ Petr Hliněný³ Daniel Lokshtanov¹ Jan Obdržálek³
Sebastian Ordyniak³ M. S. Ramanujan¹
Saket Saurabh^{1,2}

¹ University of Bergen, Bergen, Norway,
{daniello,Ramanujan.Sridharan}@ii.uib.no

² The Institute of Mathematical Sciences, Chennai, India,
saket@imsc.res.in

³ Faculty of Informatics, Masaryk University, Brno, Czech Republic.
{gajarsky,hlineny,obdrzalek,ordyniak}@fi.muni.cz

Abstract

Over the past two decades the main focus of research into first-order (FO) model checking algorithms have been sparse relational structures—culminating in the FPT-algorithm by Grohe, Kreutzer and Siebertz for FO model checking of nowhere dense classes of graphs [STOC’14], with dense structures starting to attract attention only recently. Bova, Ganian and Szeider [LICS’14] initiated the study of the complexity of FO model checking on partially ordered sets (posets). Bova, Ganian and Szeider showed that model checking *existential* FO logic is fixed-parameter tractable (FPT) on posets of bounded width, where the width of a poset is the size of the largest antichain in the poset. The existence of an FPT algorithm for general FO model checking on posets of bounded width, however, remained open. We resolve this question in the positive by giving an algorithm that takes as its input an n -element poset \mathcal{P} of width w and an FO logic formula φ , and determines whether φ holds on \mathcal{P} in time $f(\varphi, w) \cdot n^2$.

1 Introduction

Algorithmic meta-theorems are general algorithmic results applying to a whole range of problems, rather than just to a single problem alone. Such results are some of the most sought-after in algorithmic research. Many prominent algorithmic meta-theorems are about *model checking*; such theorems state that for certain kinds of logic \mathcal{L} , and all classes \mathcal{C} that have a certain structure, there is an algorithm that takes as an input a formula $\varphi \in \mathcal{L}$ and a structure $S \in \mathcal{C}$ and efficiently determines whether $S \models \varphi$. Here $S \models \varphi$ is read as “ S models φ ” or “ φ holds on S ”. Examples of theorems of this kind include the classic theorem of Courcelle [3], as well as a large body of work on model checking first-order (FO) logic [1, 4, 6, 9, 11, 13, 12, 15, 21].

Most of the research on algorithms for FO model checking has focused on graphs. On general graphs there is a naive brute-force algorithm that takes as an input an n -vertex graph G and a formula φ and determines whether $G \models \varphi$ in time $n^{O(|\varphi|)}$ by enumerating all the possible ways to instantiate the variables of φ . On the other hand, the problem is PSPACE-complete (see e.g. [7]) and encodes the CLIQUE problem, thus it admits no algorithm with running time $f(\varphi)n^{o(|\varphi|)}$ for any function f [18], assuming the Exponential Time Hypothesis [16] (ETH). Thus, assuming the ETH the naive algorithm is the best possible, up to constants in the exponent. Furthermore, FO

model checking remains PSPACE-complete on any fixed graph containing at least two vertices (again, see [7]). Hence, it is futile to look for restricted classes of graphs in which FO model checking can be done in polynomial time without restricting φ . Therefore, research has focused on obtaining algorithms with running time $f(\varphi)n^{O(1)}$ on restricted classes of graphs and other structures. Algorithms with such a running time are said to be *fixed parameter tractable* (FPT) parameterized by φ . Even though FPT algorithms are not polynomial time algorithms, due to unlimited f , they significantly outperform brute-force.

The parameterized complexity of FO model checking on *sparse* graph classes is now well understood. In 1994 Seese [21] showed an FPT algorithm for FO model checking on graphs of bounded degree. Seese’s algorithm was followed by a long line of work [11, 9, 4, 6] giving FPT algorithms for progressively larger classes of sparse graphs, culminating in the FPT algorithm of Grohe, Kreutzer and Siebertz [15] on any nowhere dense graph classes. To complement this, Kreutzer [17] and Dvořák et al. [6] proved that if a class \mathcal{C} closed under taking subgraphs is not nowhere dense, then deciding first-order properties of graphs in \mathcal{C} is not fixed-parameter tractable unless $\text{FPT}=\text{W}[1]$ (a complexity-theoretic collapse which is considered to be unlikely). Hence, this means that the result of Grohe, Kreutzer and Siebertz [15] captures *all* subgraph-closed sparse graph classes on which FO model checking is fixed parameter tractable.

However, for other types of structures, such as dense graphs or algebraic structures, the parameterized complexity of FO model checking is largely uncharted territory. Grohe [14] notes that “*it would also be very interesting to study the complexity of model-checking problems on finite algebraic structures such as groups, rings, fields, lattices, et cetera*”. From this perspective it is particularly interesting to investigate model checking problems on *partially ordered sets* (posets), since posets can be seen both as dense graphs and as algebraic structures. Motivated by Grohe’s survey [14], Bova, Ganian and Szeider [1, 2] initiated the study of FO model checking on posets. As a preliminary result, they show that FO model checking on posets parameterized by φ is not fixed parameter tractable unless $\text{FPT} = \text{W}[1]$, motivating the study of FO model checking on restricted classes of posets. Bova, Ganian and Szeider [1] identified posets of bounded width as a particularly interesting class to investigate. Their main technical contribution has been an FPT algorithm for model checking *existential* FO logic on posets of bounded width, and they left the existence of an FPT algorithm for model checking FO logic on posets of bounded width as an open problem. In subsequent work Gajarský et al. [12] gave a simpler and faster algorithm for model checking existential FO logic on posets of bounded width. Nevertheless, the existence of an FPT algorithm for model checking full FO logic remained open.

Our contribution. In this paper we resolve the open problem of Bova, Ganian and Szeider by designing a new algorithm for model checking FO logic on posets. The running time of our algorithm on an n -element poset of width w is $f(\varphi, w) \cdot n^2$. Thus our algorithm is not only FPT when parameterized by φ on posets of bounded width, it is also FPT by the compound parameter $\varphi+w$. We demonstrate the generality and applicability of our main result by showing that a simple FO-interpretation can be used to obtain an FPT-algorithm for another natural dense graph class, namely k -fold proper interval graphs. This generalizes and simplifies the main result of Ganian et al. [13].

Our algorithm is based on a new locality lemma for posets. More concretely, we show that for every poset \mathcal{P} and formula φ one can efficiently iteratively construct a directed graph D such that (a) the vertex set of D are the elements of \mathcal{P} , (b) every element of \mathcal{P} has bounded out-degree in D , and (c) it is possible to determine whether $\mathcal{P} \models \varphi$ by checking whether φ holds on sub-posets of \mathcal{P} induced by constant-radius balls in D .

The statement of our lemma sounds very similar to that of Gaifman’s locality theorem, the crucial differences being that the digraph D is not the Gaifman graph of \mathcal{P} and that D depends on the quantifier rank of φ . Indeed, constant radius balls in the Gaifman graph of constant

width posets typically contain the entire poset. Thus a naive application of Gaifman’s theorem would reduce the problem of deciding whether \mathcal{P} is a model of φ to itself. The crucial difficulty we have to overcome is that we have to make the digraph D “dense enough” so that (c) holds, while keeping it “sparse enough” so that the vertices in D still have bounded out-degree. The latter is necessary to ensure that constant radius balls in D have constant size, making it feasible to use the naive model checking algorithm for determining whether φ holds on sub-posets of \mathcal{P} induced by constant-radius balls in D . The construction of the graph D and the proof that it indeed has the desired properties relies on a delicate inductive argument thoroughly exploiting properties of posets of bounded width.

Organization of the paper. In Section 2 we set up the definitions and the necessary notation. In Section 3 we define the digraph D used in our poset locality lemma, and prove some useful structural properties of D . In Section 4 we prove the locality lemma for posets and give the FPT algorithm for FO logic model checking. We then proceed to show in Section 5 how our algorithm can be used to give an FPT algorithm for model checking FO logic on k -fold proper interval graphs. Finally, in Section 6 we conclude with a discussion of further research directions.

2 Preliminaries

2.1 Graphs and Posets

We deal with directed graphs (shortly *digraphs*) whose vertices and arcs bear auxiliary labels, and which may contain parallel arcs. For a directed graph D , a vertex $v \in V(D)$, and an integer r , we denote by $R_r^D(v)$, the set of vertices of D that are reachable from v via a directed path of length at most r . Slightly abusing the notation, we extend the function R_r^D to multiple vertices as follows; $R_r^D(v_1, \dots, v_k) = \bigcup_{i=1}^k R_r^D(v_i)$. Moreover, for $v, v' \in V(D)$ we denote by $\text{dist}_D(v, v')$ the length of a shortest directed path from v to v' in D .

A *poset* \mathcal{P} is a pair $(P, \leq^{\mathcal{P}})$ where P is a finite set and $\leq^{\mathcal{P}}$ is a reflexive, anti-symmetric, and transitive binary relation over P . The *size* of a poset $\mathcal{P} = (P, \leq^{\mathcal{P}})$ is $\|\mathcal{P}\| := |P|$. We say that p and p' are *incomparable* (in \mathcal{P}), denoted $p \parallel^{\mathcal{P}} p'$, if neither $p \leq^{\mathcal{P}} p'$ nor $p' \leq^{\mathcal{P}} p$ hold. We say that p' is *above* p (and p is *below* p') if $p \leq^{\mathcal{P}} p'$ and $p \neq p'$. A *chain* C of \mathcal{P} is a subset of P such that $x \leq^{\mathcal{P}} y$ or $y \leq^{\mathcal{P}} x$ for every $x, y \in C$. A *chain partition* of \mathcal{P} is a tuple (C_1, \dots, C_k) such that $\{C_1, \dots, C_k\}$ is a partition of P and for every i with $1 \leq i \leq k$ the poset induced by C_i is a chain of \mathcal{P} . An *anti-chain* A of \mathcal{P} is a subset of P such that for all $x, y \in A$ it is true that $x \parallel^{\mathcal{P}} y$. The *width* of a poset \mathcal{P} , denoted by $\text{width}(\mathcal{P})$ is the maximum cardinality of any anti-chain of \mathcal{P} .

Proposition 2.1 ([8, Theorem 1.]). *Let \mathcal{P} be a poset. Then in time $\mathcal{O}(\text{width}(\mathcal{P}) \cdot \|\mathcal{P}\|^2)$, it is possible to compute both $\text{width}(\mathcal{P}) = w$ and a corresponding chain partition (C_1, \dots, C_w) of \mathcal{P} .*

2.2 Parameterized Complexity

Here we introduce the most basic concepts of parameterized complexity theory. For more details, we refer to the many existing text books on the topic [5, 10, 19]. An instance of a parameterized problem is a pair $\langle x, k \rangle$ where x is the input and k a parameter. A parameterized problem \mathcal{P} is *fixed-parameter tractable (FPT)* if, for every instance $\langle x, k \rangle$, it can be decided whether $\langle x, k \rangle \in \mathcal{P}$ in time $f(k) \cdot |x|^c$, where f is a computable function, and c is a constant.

2.3 First-order Logic

In this paper we deal with the, well known, relational first-order (FO) logic. Formulas of this logic are built from (a finite set of) variables, relational symbols, logical connectives (\wedge, \vee, \neg)

and quantifiers (\exists, \forall). A *sentence* is a formula with no free variables. We restrict ourselves to formulas that are in *negation normal form*; a first-order formula is in negation normal form if all negation symbols occur only in front of the atoms. Obviously, any first-order formula can be, in a linear time, converted into an equivalent one in negation normal form.

The problem we are interested in is a *model checking problem* for FO formulas on posets, which is formally defined as follows:

POSET FO MODEL CHECKING Input: A first-order sentence φ and a poset $\mathcal{P} = (P, \leq^{\mathcal{P}})$. Question: Is it true $\mathcal{P} \models \varphi$, i.e., is \mathcal{P} a model of φ ?	Parameter: $\text{width}(\mathcal{P}), \varphi $
--	--

All first-order formulas in this paper are evaluated over posets as follows. The vocabulary consists of the one binary relation $\leq^{\mathcal{P}}$ and a finite set of arbitrary unary relations (“colors” of poset elements). Atoms of these FO formulas can be equalities between variables ($x = y$), applications of the predicate $\leq^{\mathcal{P}}$ (with the natural meaning of $x \leq^{\mathcal{P}} y$ in the poset \mathcal{P}), or applications of one of the unary predicates $c(x)$ (with the meaning that x is of color c). For a more detailed treatment of the employed setting, we refer the reader to [1].

While for any fixed sentence φ one can easily decide whether $\mathcal{P} \models \varphi$ in polynomial time, by a brute-force expansion of all the quantifiers, such a solution is not FPT since the exponent depends on φ . Our aim is to provide an FPT solution in the case when we additionally parameterize the input by the width of \mathcal{P} .

It is well known that the model checking problem for almost any logic can be formulated as finding a winning strategy in an appropriate model checking game, often called *Hintikka game* (see e.g. [7]). In our case the game $\mathcal{G}(\mathcal{P}, \varphi)$ for a poset \mathcal{P} and an FO formula $\varphi = \varphi(x_1, \dots, x_k)$ in negation normal form (where x_1, \dots, x_k are the free variables of φ) is defined as follows:

The game is played by two players, the existential player (Player \exists , Verifier), who tries to prove that $\mathcal{P} \models \varphi$, and the universal player (Player \forall , Falsifier), who tries to disprove that claim. The positions of this game $\mathcal{G}(\mathcal{P}, \varphi)$ are pairs (ψ, β) , where $\psi \equiv \psi(x_1, \dots, x_\ell)$ is a subformula of φ , and $\beta : \{x_1, \dots, x_\ell\} \rightarrow P$ assigns free variables of ψ elements of the poset \mathcal{P} . We write $\psi(p_1, \dots, p_\ell)$ for a position (ψ, β) , where $\beta(x_i) = p_i$ for the free variables x_i of ψ . The initial position of the game is (φ, β_0) , where β_0 is the initial assignment (if φ has free variables, or β_0 is empty).

The game is played as follows: the existential player (Verifier) moves from positions associated with disjunctions and formulas starting with the existential quantifier. From a position $\psi_1 \vee \psi_2$ he moves to either ψ_1 or ψ_2 . From a position $\psi(p_1, \dots, p_i) \equiv \exists y. \psi'(p_1, \dots, p_i, y)$ he moves to any position $\psi'(p_1, \dots, p_i, p)$, where $p \in P$. The universal player plays similarly from conjunctions and universally quantified formulas. At atoms which, in our case, are the positions $\sigma(p_1)$ of the form $c(p_1)$, or $\sigma(p_1, p_2)$ of the form $p_1 = p_2, \neg(p_1 = p_2), p_1 \leq^{\mathcal{P}} p_2$ or $\neg(p_1 \leq^{\mathcal{P}} p_2)$, the existential player wins if $\mathcal{P} \models \sigma(p_1, p_2)$, and otherwise the universal player wins. The equivalence of these games to the standard semantic of FO is given by the following claim:

Proposition 2.2. *The existential player has a winning strategy in the Hintikka game $\mathcal{G}(\mathcal{P}, \varphi)$ for a poset \mathcal{P} and a first-order sentence φ if, and only if, $\mathcal{P} \models \varphi$.*

3 Poset Structure and Types

For the rest of the paper we fix a poset $\mathcal{P} = (P, \leq^{\mathcal{P}})$ of width w , a mapping $\lambda : P \rightarrow \Lambda$ where Λ is a (fixed) finite set of *colors*, and a chain partition (C_1, \dots, C_w) of \mathcal{P} . To emphasize that \mathcal{P} is associated with an auxiliary coloring λ , we sometimes call \mathcal{P} a *colored poset*. For $p \in C_j$ with $1 \leq j \leq w$, we denote by $C(p)$ the chain C_j . The purpose of this section is to find a description

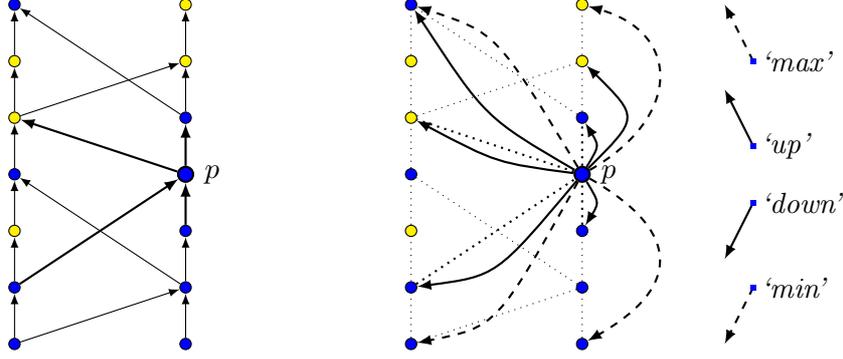


Figure 1: The picture, on the left, shows an upward-directed Hasse diagram of a bicolored poset \mathcal{P} (where \mathcal{P} is the reflexive and transitive closure of it). On the right, the picture shows the arcs of D_0 starting from a selected element $p \in P$, as by Definition 3.1.

of the poset \mathcal{P} structure suitable for applying “locality tools” of finite model theory. This turns out to be a delicate job requiring a careful inductive definition.

For an integer $s \geq 0$, we set $r_s := 3 \cdot 4^s - 1$ and inductively define a) a labeling function $\tau_s : P \rightarrow \mathbb{N}$, and b) a vertex-labeled and arc-labeled directed graph D_s on the vertex set $V(D_s) := P$ as follows:

Definition 3.1. For an integer $s \geq 0$ and an element $p \in P$, we shortly denote by $P_s(p)$ the set $R_{r_s}^{D_s}(p)$ (i.e., the set of vertices reachable in D_s from p at distance $\leq r_s$). We let $\tau_0(p) := \langle \lambda(p), j \rangle$, where j is the index s.t. $C_j = C(p)$. Inductively for every integer $s \geq 0$, we define D_s as the digraph with the vertex set P and vertex labels given by τ_s , containing the following arcs:

- for every $p \in P$ and every $j \in \{1, \dots, w\}$, D_s contains an arc from p with label ‘max’ to the topmost element of C_j ;
- for every $p \in P$ and every $j \in \{1, \dots, w\}$, D_s contains an arc from p with label ‘min’ to the bottommost element of C_j ;
- for every $p \in P$, every $j \in \{1, \dots, w\}$, and every $t \in \{\tau_s(q) \mid q \in C_j\}$, D_s contains an arc from p to p' with label ‘up’, where $p' \neq p$ is the bottommost element of C_j such that $\tau_s(p') = t$ and $p \leq^{\mathcal{P}} p'$ (if such an element p' exists);
- for every $p \in P$, every $j \in \{1, \dots, w\}$, and every $t \in \{\tau_s(q) \mid q \in C_j\}$, D_s contains an arc from p to p' with label ‘down’, where $p' \neq p$ is the topmost element of C_j such that $\tau_s(p') = t$ and $p' \leq^{\mathcal{P}} p$ (if such an element p' exists).

Then, having defined D_s as above, we set for every element $p \in P$

- $\tau_{s+1}(p)$ to be the isomorphism type of a relational structure $\mathcal{A}_s^{\mathcal{P}}(p)$, where $\mathcal{A}_s^{\mathcal{P}}(p)$ is formed by the vertex- and arc-labeled induced subdigraph $D_s[P_s(p)]$ rooted at p with the additional binary relation $\leq^{\mathcal{P}}$ restricted to $P_s(p)$.

The values $\tau_s(p)$, $p \in P$, will also be called the *types of rank s* (of elements of \mathcal{P}), where the rank will often be implicit from the context. It is useful to notice that the considered coloring λ of the poset \mathcal{P} elements is fully determined by their types of rank 0 in D_0 (and so also by their types of any higher rank). Therefore, we may skip an explicit reference to λ in the rest of this section.

The definition of D_0 is illustrated in Figure 1. Informally, the type of an element $p \in P$ captures its “local neighborhood” (which is growing in size with the rank s), and the digraph D_s contains ‘up’-arcs (‘down’-arcs) from p to the next higher (next lower) elements of \mathcal{P} of

each appearing type. Moreover, there are shortcut arcs, labeled ‘*min*’ and ‘*max*’, from p to the extreme elements of each chain of \mathcal{P} . It is important that, since we use a fixed finite number of colors in \mathcal{P} and since \mathcal{P} is of bounded width, the outdegrees in D_s are inductively bounded for every fixed s independently of the size of \mathcal{P} .

For start we need the following basic properties of the digraph D_s and the labeling function τ_s , which are easy to prove. The first two of these simple claims establish that the sequence of labeled digraphs D_0, D_1, D_2, \dots indeed presents an increasingly finer resolution of a “local structure” of the poset \mathcal{P} . For all the claims, let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a poset and $P_s, D_s, \mathcal{A}_s^{\mathcal{P}}$ and τ_s be as in Definition 3.1.

Lemma 3.2. *For every $p, p' \in P$ and $s \geq 0$, if $\tau_s(p) \neq \tau_s(p')$, then also $\tau_{s+1}(p) \neq \tau_{s+1}(p')$.*

Proof. Since $\tau_s(p) \neq \tau_s(p')$, it follows that p and p' have different vertex labels in D_s . Hence, the isomorphism types of $\mathcal{A}_s^{\mathcal{P}}(p)$ and $\mathcal{A}_s^{\mathcal{P}}(p')$ are not the same either and hence $\tau_{s+1}(p) \neq \tau_{s+1}(p')$. \square

Lemma 3.3. *For every $s \geq 0$, if D_s contains an arc from some vertex $p \in P$ to some $p' \in P$, then D_{s+1} also contains an arc from p to p' with the same label. In other words, D_s is a spanning subdigraph of D_{s+1} (neglecting the vertex-labels).*

Proof. Let $p, p' \in P$ and assume there is an arc a from p to p' in D_s . We need to show that there is also an arc from p to p' in D_{s+1} that has the same label as a . Depending on the arc-label of a in D_s we distinguish the following cases:

1. Assume that the label of a is ‘*up*’. Then, p' is the bottommost element of type $\tau_s(p')$ on the chain $C(p')$ such that $p \leq^{\mathcal{P}} p'$. Hence, for any element $p'' \in C(p')$ such that $p \leq^{\mathcal{P}} p'' \leq^{\mathcal{P}} p'$, it holds that $\tau_s(p'') \neq \tau_s(p')$. By Lemma 3.2, we obtain that also $\tau_{s+1}(p'') \neq \tau_{s+1}(p')$. Consequently, p' is also the bottommost element of type $\tau_{s+1}(p')$ on the chain $C(p')$ with $p \leq^{\mathcal{P}} p'$. This shows that D_{s+1} contains an arc from p to p' with label ‘*up*’, as required.
2. The argument for the case of label ‘*down*’ is similar to the previous case.
3. The claim trivially holds for the labels ‘*min*’ and ‘*max*’.

\square

Another simple property of Definition 3.1 is that pairs of arcs of the same vertex- and arc-labels in the digraph D_s never “cross one another”, which is formalized as follows:

Lemma 3.4. *Let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a poset and P_s, D_s and τ_s be as in Definition 3.1. Assume that $p, p', q, q' \in P$ are such that $\tau_s(p) = \tau_s(p')$ and $\tau_s(q) = \tau_s(q')$, and that both (p, q) and (p', q') are arcs of the same label in D_s . If $p \leq^{\mathcal{P}} p'$ then $q \leq^{\mathcal{P}} q'$.*

Proof. Note that $p' \in C(p)$ and $q' \in C(q)$, and so the pairs are comparable. If the label of (p, q) is ‘*min*’ or ‘*max*’, then $q = q'$ and the claim holds. Assume that the label of (p, q) is ‘*up*’ and that $q' \leq^{\mathcal{P}} q$. Then $p \leq^{\mathcal{P}} p' \leq^{\mathcal{P}} q'$ and so, by Definition 3.1, the arc (p, q) should point from p to q' in D_s , a contradiction. Therefore, $q \leq^{\mathcal{P}} q'$ as required. The case of label ‘*down*’ is similar. \square

The subsequent claims are more involved and technical. Informally, they together show that for any $p_1, \dots, p_k \in P$, a property or relation of other element(s) of \mathcal{P} to p_1, \dots, p_k can also be observed in a given bounded neighborhood of p_1, \dots, p_k in D_s . Importantly, the richer local property is observed, the higher index s in D_s is used. The easy base case of $s = 0$ is covered by Lemma 3.5 while the general case of s is inductively established by Lemma 3.6 and reformulated in Corollary 3.7. We refer to Section 4 for details on using these claims.

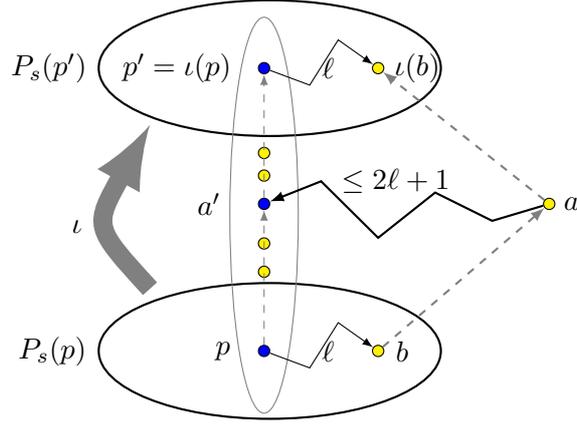


Figure 2: An illustration of the statement of Lemma 3.6: the dashed arcs depict the poset relation $\leq^{\mathcal{P}}$ while the solid arcs represent directed paths in the digraph D_s .

Lemma 3.5. *Let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a poset and D_0 the digraph defined in Definition 3.1. For any $k \geq 1$ and $p, p_1, \dots, p_k \in P$, there exists an element $p' \in R_2^{D_0}(p_1, \dots, p_k)$ such that $\tau_0(p') = \tau_0(p)$, and p', p are in the same relation with respect to all of p_1, \dots, p_k in \mathcal{P} : formally, for every $i \in \{1, \dots, k\}$, it holds that $p' \leq^{\mathcal{P}} p_i$ if and only if $p \leq^{\mathcal{P}} p_i$, and $p_i \leq^{\mathcal{P}} p'$ if and only if $p_i \leq^{\mathcal{P}} p$.*

Proof. Choose p' to be the topmost element of the chain $C(p)$ such that $\tau_0(p') = \tau_0(p)$ and, for every element $q \in C(p)$ with $p \leq^{\mathcal{P}} q \leq^{\mathcal{P}} p'$, it holds that $q \notin R_1^{D_0}(p_1, \dots, p_k)$. Then, clearly, $p \leq^{\mathcal{P}} p'$ (while it might happen $p = p'$) and $p' \in R_2^{D_0}(p_1, \dots, p_k)$ by the definition. It remains to show that p' is in the same relation with respect to p_1, \dots, p_k in \mathcal{P} as p . Assume for a contradiction that this is not the case. Then there is an index $i \in \{1, \dots, k\}$ such that p and p' are in a different relation towards p_i . Because $p \leq^{\mathcal{P}} p'$, we have that either (I) $p \leq^{\mathcal{P}} p_i$ but $p' \not\leq^{\mathcal{P}} p_i$, or (II) $p_i \leq^{\mathcal{P}} p'$ but $p_i \not\leq^{\mathcal{P}} p$.

In the case (I), either $p_i \in C(p)$ and hence p_i is in between p and p' on the chain $C(p)$ contradicting our choice of p' , or $p_i \notin C(p)$ but D_0 contains an arc with label ‘down’ from p_i to some element on the chain $C(p)$ between p and p' , which again contradicts our choice of p' . In the case (II), either $p_i \in C(p)$ and hence p_i is in between p and p' on the chain $C(p)$ contradicting our choice of p' , or $p_i \notin C(p)$ but D_0 contains an arc with label ‘up’ from p_i to some element on the chain $C(p)$ between p and p' , which again contradicts our choice of p' . \square

We use the following shorthand notation. For $p, p', q \in P$ we say that q *discerns* p from p' , with respect to $\leq^{\mathcal{P}}$, if (at least) one of the following four conditions holds true; $p \leq^{\mathcal{P}} q$ and $p' \not\leq^{\mathcal{P}} q$, $p \not\leq^{\mathcal{P}} q$ and $p' \leq^{\mathcal{P}} q$, $q \leq^{\mathcal{P}} p$ and $q \not\leq^{\mathcal{P}} p'$, or $q \not\leq^{\mathcal{P}} p$ and $q \leq^{\mathcal{P}} p'$. For example, the conclusion of Lemma 3.5 is equivalent to saying “neither of p_1, \dots, p_k discerns p from p' ”. Since the poset \mathcal{P} is fixed for this section, we will often skip an explicit reference to $\leq^{\mathcal{P}}$.

Lemma 3.6. *Let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a poset, $s \geq 0$ an integer, D_s and τ_s be as in Definition 3.1, and $p, p' \in P$. Assume that $p \leq^{\mathcal{P}} p'$ and $\tau_{s+1}(p) = \tau_{s+1}(p')$, where the latter is witnessed by an isomorphism $\iota : P_s(p) \rightarrow P_s(p')$ of the structures $\mathcal{A}_s^{\mathcal{P}}(p)$ and $\mathcal{A}_s^{\mathcal{P}}(p')$. If $a \in P$ and $b \in P_s(p)$ are such that a discerns b from $\iota(b)$ (with respect to $\leq^{\mathcal{P}}$), then there exists a directed path from a to some element $a' \in C(p)$ of type $\tau_{s+1}(p)$ with $p \leq^{\mathcal{P}} a' \leq^{\mathcal{P}} p'$, in D_s of length at most $2\ell + 1$ where $\ell = \text{dist}_{D_s}(p, b)$.*

The technical statement of Lemma 3.6 deserves an informal explanation. For start, if elements $p, p' \in P$ such that $p' \in C(p)$ are discerned in the poset \mathcal{P} by an element $a \in P$ then, clearly, $a = a' \in C(p)$ or the Hasse diagram of \mathcal{P} contains an arc between a and some $a' \in C(p)$

such that a' lies between p and p' on $C(p)$. Lemma 3.6 then largely extends this simple observation to the setting of Definition 3.1 and for discerned elements in neighborhoods of p and p' . The statement is illustrated in Figure 2.

Proof of Lemma 3.6. Notice that the assumption $\tau_{s+1}(p) = \tau_{s+1}(p')$ immediately implies $p' \in C(p)$. Let $F = (b_0 = p, b_1, \dots, b_\ell = b)$ be a shortest directed path from p to b in D_s . Then, $\iota(F) = (\iota(b_0) = p', \iota(b_1), \dots, \iota(b_\ell) = \iota(b))$ is a directed path in D_s that has the same arc-labels and vertex-labels as F . Observe that since a discerns b from $\iota(b)$, it holds that $b \neq \iota(b)$ and hence also $b_i \neq \iota(b_i)$ for every $i \in \{0, \dots, \ell\}$. Moreover, $p \leq^{\mathcal{P}} p' = \iota(p)$ implies $b \leq^{\mathcal{P}} \iota(b)$ by a simple inductive argument using Lemma 3.4 along the path F .

We will show the existence of $a' \in C(p)$ and of the required path from a to a' in D_s by proving the following slightly stronger statement:

- For every $i \in \{0, \dots, \ell\}$, there is an element $a_i \in C(b_i)$ of type $\tau_{s+1}(b_i)$ such that $b_i \leq^{\mathcal{P}} a_i \leq^{\mathcal{P}} \iota(b_i)$ and D_s contains a directed path from a to a_i of length at most $2(\ell - i) + 1$.

We prove the statement via induction on $(\ell - i)$ starting from $i = \ell$.

For $i = \ell$, we obtain a_ℓ as follows. If $a \in C(b)$, then $b \leq^{\mathcal{P}} a \leq^{\mathcal{P}} \iota(b)$ because a discerns b from $\iota(b)$, and let $a_\ell \in C(b)$ be the topmost element of type $\tau_{s+1}(b)$ such that $a_\ell \leq^{\mathcal{P}} a$. Clearly, $b \leq^{\mathcal{P}} a_\ell \leq^{\mathcal{P}} \iota(b)$ and D_s contains an arc with label ‘down’ from a to a_ℓ , as required. So assume that $a \notin C(b)$. Then, again, $b \leq^{\mathcal{P}} \iota(b)$ and a discerns b from $\iota(b)$, and so at least one of $b \leq^{\mathcal{P}} a$ or $a \leq^{\mathcal{P}} \iota(b)$ holds true.

In the former case, let $a_\ell \in C(b)$ be the topmost element of type $\tau_{s+1}(b)$ such that $a_\ell \leq^{\mathcal{P}} a$. Since $b \leq^{\mathcal{P}} a$ but $\iota(b) \not\leq^{\mathcal{P}} a$, it follows that $b \leq^{\mathcal{P}} a_\ell \leq^{\mathcal{P}} \iota(b)$. Furthermore, D_s contains an arc with label ‘down’ from a to a_ℓ , as required. In the later case, let $a_\ell \in C(b)$ be the bottommost element of type $\tau_{s+1}(b)$ such that $a \leq^{\mathcal{P}} a_\ell$. Since $a \leq^{\mathcal{P}} \iota(b)$ but $a \not\leq^{\mathcal{P}} b$, it again follows that $b \leq^{\mathcal{P}} a_\ell \leq^{\mathcal{P}} \iota(b)$. Furthermore, D_s contains an arc with label ‘up’ from a to a_ℓ , as required.

Now, consider $i \in \{0, \dots, \ell - 1\}$ and that we have already shown the claim for $i + 1$. Hence, by the induction hypothesis, we can assume that there exists $a_{i+1} \in C(b_{i+1})$ of type $\tau_{s+1}(b_{i+1})$ with $b_{i+1} \leq^{\mathcal{P}} a_{i+1} \leq^{\mathcal{P}} \iota(b_{i+1})$ such that D_s contains a directed path from a to a_{i+1} of length at most $2(\ell - i - 1) + 1$.

We need to show that there is $a_i \in C(b_i)$ of type $\tau_{s+1}(b_i)$ with $b_i \leq^{\mathcal{P}} a_i \leq^{\mathcal{P}} \iota(b_i)$ such that D_s contains a directed path of length at most 2 from a_{i+1} to a_i . We distinguish the following cases, depending on the arc-label of the arc (b_i, b_{i+1}) in D_s .

1. Assume that the label of (b_i, b_{i+1}) is ‘up’. See Figure 3.

If $\iota(b_i) \not\leq^{\mathcal{P}} a_{i+1}$, then let $a_i \in C(b_i)$ be the topmost element in $C(b_i)$ of type $\tau_{s+1}(b_i)$ such that $a_i \leq^{\mathcal{P}} a_{i+1}$. Then, D_s contains an arc with label ‘down’ from a_{i+1} to a_i (or, possibly, $a_i = a_{i+1}$). Since $\iota(b_i) \leq^{\mathcal{P}} a_i$ would contradict $\iota(b_i) \not\leq^{\mathcal{P}} a_{i+1}$, we have $a_i \leq^{\mathcal{P}} \iota(b_i)$. Since $b_i \leq^{\mathcal{P}} b_{i+1} \leq^{\mathcal{P}} a_{i+1}$ and a_i is the topmost one of its kind in $C(b_i)$, it holds that $b_i \leq^{\mathcal{P}} a_i$, as required.

Otherwise, $\iota(b_i) \leq^{\mathcal{P}} a_{i+1}$ and because the label of the arc $(\iota(b_i), \iota(b_{i+1}))$ in D_s is ‘up’, we have $a_{i+1} = \iota(b_{i+1})$. Let h be the topmost element in $C(b_{i+1})$ of type $\tau_{s+1}(b_{i+1})$ with $h \leq^{\mathcal{P}} a_{i+1}$. Then, D_s contains an arc with label ‘down’ from a_{i+1} to h and $b_{i+1} \leq^{\mathcal{P}} h$ (recall $b_{i+1} \neq \iota(b_{i+1})$). Because of (the existence of) the arc $(\iota(b_i), \iota(b_{i+1}))$ in D_s , we obtain that $\iota(b_i) \not\leq^{\mathcal{P}} h$. Let $a_i \in C(b_i)$ be the topmost element in $C(b_i)$ of type $\tau_{s+1}(b_i)$ such that $a_i \leq^{\mathcal{P}} h$. Then, D_s contains an arc with label ‘down’ from h to a_i and hence a directed path from a_{i+1} to a_i of length 2 (or, possibly, $a_i = h$). Clearly, $a_i \leq^{\mathcal{P}} \iota(b_i)$. Since $b_i \leq^{\mathcal{P}} b_{i+1} \leq^{\mathcal{P}} h$ and a_i is the topmost one of this kind in $C(b_i)$, it holds that $b_i \leq^{\mathcal{P}} a_i$, as required.

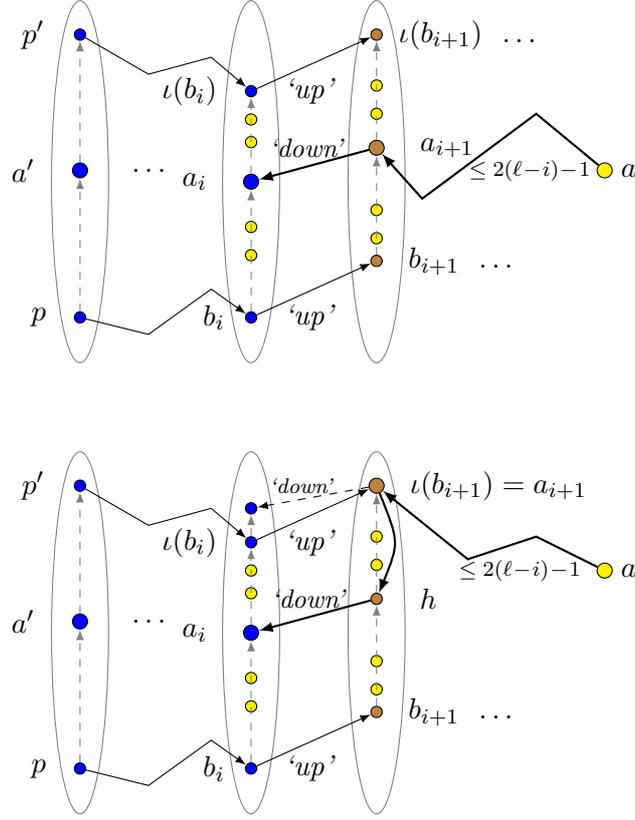


Figure 3: An illustration of the proof of Lemma 3.6: top – the inductive step in the case of $\iota(b_i) \not\leq^P a_{i+1}$; bottom – the inductive step in the case of $\iota(b_i) \leq^P a_{i+1}$.

2. Assume that the label of (b_i, b_{i+1}) is ‘down’. The situation is completely symmetric to the previous case, and the same arguments prove the conclusion (just exchanging \leq^P with \geq^P and ‘down’ with ‘up’).
3. Assume that the label of (b_i, b_{i+1}) is ‘max’ or ‘min’. That would imply $b_{i+1} = \iota(b_{i+1})$, a contradiction. This case cannot happen. \square

The way we shall use Lemma 3.6 in Section 4 can be informally summarized as follows. If p and p' are elements of the same type which are next to each other on their chain in \mathcal{P} , then their D_s -neighborhoods appear “the same” with respect to \leq^P to all poset elements which are sufficiently far away from p as measured by D_{s+1} . The precise formulation is next.

Corollary 3.7. *Let $\mathcal{P} = (P, \leq^P)$ be a poset, $s \geq 0$ an integer, the digraphs D_s and D_{s+1} be as in Definition 3.1, $p \neq p' \in P$ be such that $\tau_{s+1}(p) = \tau_{s+1}(p')$, and the isomorphism map $\iota : P_s(p) \rightarrow P_s(p')$ be as in Lemma 3.6. Assume, moreover, that D_{s+1} contains an arc (p', p) with label ‘down’. For any given $p_1, \dots, p_k \in P$, $k \geq 1$, denote by $S := R_{r_s}^{D_s}(p_1, \dots, p_k)$ and define a mapping $f : S \cup P_s(p) \rightarrow S \cup P_s(p')$ such that $f(e) = \iota(e)$ for $e \in P_s(p)$ and $f(e) = e$ otherwise. If $p \notin R_{r_{s+1}-r_s}^{D_{s+1}}(p_1, \dots, p_k)$, then f is a color-preserving isomorphism between the induced subposets $\mathcal{P}[S \cup P_s(p)]$ and $\mathcal{P}[S \cup P_s(p')]$.*

The statement of Corollary 3.7 is illustrated in Figure 4.

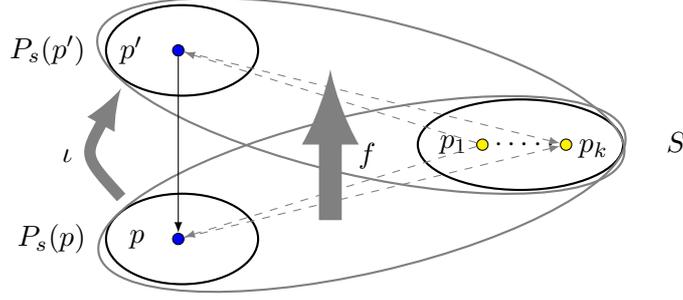


Figure 4: An illustration of the statement of Corollary 3.7: f is a color-preserving isomorphism between the induced subposets $\mathcal{P}[S \cup P_s(p)]$ and $\mathcal{P}[S \cup P_s(p')]$.

Proof. For the sake of contradiction assume that $p \notin R_{r_{s+1}-r_s}^{D_{s+1}}(p_1, \dots, p_k)$ and f is not a poset isomorphism. Then either f is not a bijection or there are two elements $a, b \in S \cup P_s(p)$ such that a and b are in a different relation with respect to $\leq^{\mathcal{P}}$ than $f(a)$ and $f(b)$.

Observe that in the latter case, since a and b are in a different relation with respect to $\leq^{\mathcal{P}}$ than $f(a)$ and $f(b)$, and by the definition of f we clearly see that, up to symmetry, $a \in S \setminus P_s(p)$ and $b \in P_s(p)$. Hence, the latter case means that $a = f(a)$ discerns b from $f(b) = \iota(b)$.

In the former case consider the function f' defined analogously to the function f but from the ‘‘perspective of p' ’’, i.e., $f' : S \cup P_s(p') \rightarrow S \cup P_s(p)$ is defined by setting $f'(e) = \iota^{-1}(e)$ for every $e \in P_s(p')$ and $f'(e) = e$ otherwise. Then f is a bijection if and only if f' is. Moreover, f is a bijection if and only if both f and f' are injective. If f is not injective, then there exists a $b \in P_s(p)$ such that $a = f(b) = \iota(b) \in S \setminus P_s(p)$. Since $a \neq b$, the element a discerns b from $\iota(b) = a$. On the other hand, if f' is not injective, then there exists a $b' \in P_s(p')$ such that $a' = f'(b') = \iota^{-1}(b') \in S \setminus P_s(p')$. Since $a' \neq b'$, the element a' discerns a' from b' . Setting $a := a'$ and $b := a'$, we again see that a discerns $b = a'$ from $\iota(b) = b'$.

In either of the three subcases above, the elements a and b satisfy the conditions of Lemma 3.6. Hence there exists a directed path from a to an element $a' \in C(p)$ of type $\tau_{s+1}(p)$ with $p \leq^{\mathcal{P}} a' \leq^{\mathcal{P}} p'$ in D_s of length at most $2r_s + 1$. Because of Lemma 3.3 this path exists also in D_{s+1} . Furthermore, since (p', p) is an arc with label ‘‘down’’ in D_{s+1} , there are no elements of type $\tau_{s+1}(p)$ on $C(p)$ ‘‘between’’ p and p' , and so either $a' = p$ or $a' = p'$. Hence, there is a directed path from a to p in D_{s+1} of length at most $2r_s + 1 + 1$.

Finally, since $a \in S$, there is a directed path from some of the elements p_1, \dots, p_k to a of length at most r_s . In a summary, there exists a directed path in D_{s+1} from one of p_1, \dots, p_k to p of length at most (recall $r_s = 3 \cdot 4^s - 1$)

$$2r_s + 2 + r_s = (4r_s + 3) - r_s - 1 = r_{s+1} - r_s - 1$$

which contradicts the assumption $p \notin R_{r_{s+1}-r_s}^{D_{s+1}}(p_1, \dots, p_k)$.

Color-preservation by f immediately follows from the same property of ι . \square

4 The Model Checking Algorithm

We would like to use the Hintikka game to solve the poset FO model checking problem $\mathcal{P} \models \varphi$, by Proposition 2.2. Though, the number of possible distinct plays in the game $\mathcal{G}(\mathcal{P}, \varphi)$ grows roughly as $\mathcal{O}(\|\mathcal{P}\|^{|\varphi|})$ which is not FPT. To resolve this problem, we are going to show that in fact only a small subset of all plays of the game $\mathcal{G}(\mathcal{P}, \varphi)$ is necessary to determine the outcome—only those plays which are suitably locally constrained with the use of Definition 3.1.

Let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a poset, with an implicitly associated coloring $\lambda : P \rightarrow \Lambda$ as in Section 3, and φ an FO formula. The r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \varphi)$ (where “ r ” refers to the sequence $r_s = 3 \cdot 4^s - 1$, cf. Definition 3.1) is played on the same set of positions by the same rules as the ordinary Hintikka game $\mathcal{G}(\mathcal{P}, \varphi)$, with the following additional restriction: for each $Q \in \{\exists, \forall\}$, and for any position of the form $\psi(p_1, \dots, p_i) \equiv Qy. \psi'(p_1, \dots, p_i, y)$ where $i \geq 1$, Player Q has to move to a position $\psi'(p_1, \dots, p_i, p)$ such that $p \in R_{r_q - r_{q-1}}^{D_q}(p_1, \dots, p_i)$, where $q \geq 1$ is the quantifier rank of ψ' . If $q = 0$, i.e., for quantifier-free ψ' , the restriction is $p \in R_{r_0}^{D_0}(p_1, \dots, p_i)$.

Lemma 4.1. *Let \mathcal{P} be a poset, $p_1, \dots, p_j \in P$, and φ be an FO formula of quantifier rank q . In the r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \varphi)$ with an initial position $\varphi(p_1, \dots, p_j)$, $j \geq 1$, every reachable game position $\psi(p_1, \dots, p_k)$, $k > j$, is such that*

$$p_{j+1}, \dots, p_k \in R_{r_{q-1}}^{D_{q-1}}(p_1, \dots, p_j).$$

Proof. We proceed by induction on $(k - j)$, starting with the trivial degenerate case $k = j$ and proving a stronger statement

$$p_{j+1}, \dots, p_k \in R_{r_{q-1} - r_{q+j-k-1}}^{D_{q-1}}(p_1, \dots, p_j).$$

Let the claim hold for a position $\psi(p_1, \dots, p_{k-1})$ and consider a next position $\psi'(p_1, \dots, p_{k-1}, p_k)$, where the quantifier rank of ψ' is $q' = q - (k - j)$. Then, by the definition, p_k is at distance at most $r_{q'} - r_{q'-1}$ in $D_{q'}$ from one of the elements p_1, \dots, p_{k-1} , and the same holds also in D_{q-1} by Lemma 3.3 as $q' \leq q - 1$. Since each of p_1, \dots, p_{k-1} is at distance at most $r_{q-1} - r_{q+j-(k-1)-1} = r_{q-1} - r_{q'}$ from one of p_1, \dots, p_j (and this holds also for $k - 1 = j$), we get that p_k is at distance at most $r_{q-1} - r_{q'} + r_{q'} - r_{q'-1} = r_{q-1} - r_{q'-1} = r_{q-1} - r_{q+j-k-1}$ from one of p_1, \dots, p_j in D_{q-1} . \square

Now we get to the crucial technical claim of this paper. For $Q \in \{\exists, \forall\}$, and with a neglectable abuse of terminology, we say that Player Q *wins* the (ordinary or r -local) Hintikka game *from a position* $\psi(p_1, \dots, p_i)$ if Player Q has a winning strategy in the game $\mathcal{G}(\mathcal{P}, \psi)$ (or in $\mathcal{G}_r(\mathcal{P}, \psi)$, respectively) with the initial position $\psi(p_1, \dots, p_i)$. Otherwise, Player Q *loses* the game.

Lemma 4.2. *Let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a colored poset and φ be an FO formula in negation normal form. For $Q \in \{\exists, \forall\}$ and $i \geq 0$, consider a position $\psi(p_1, \dots, p_i)$ in the Hintikka game $\mathcal{G}(\mathcal{P}, \varphi)$. If Player Q wins the game $\mathcal{G}(\mathcal{P}, \varphi)$ from the position $\psi(p_1, \dots, p_i)$, then Player Q wins also the r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \psi)$ from an initial position $\psi(p_1, \dots, p_i)$.*

Proof. Let \bar{Q} denote the other player, that is, $\{Q, \bar{Q}\} = \{\exists, \forall\}$. For the sake of contradiction, assume that we have got a counterexample with \mathcal{P} , φ , and $\psi(p_1, \dots, p_i)$ of quantifier rank $q + 1$; meaning that Player Q wins $\mathcal{G}(\mathcal{P}, \varphi)$ from $\psi(p_1, \dots, p_i)$ but Q loses $\mathcal{G}_r(\mathcal{P}, \psi)$ from initial $\psi(p_1, \dots, p_i)$. Assume, moreover, that $Q \in \{\exists, \forall\}$ and the counterexample are chosen such that the pair $\langle q, |\psi| \rangle$ is lexicographically minimal.

Since a game move associated with a conjunction or disjunction of formulas, and a (possible) initial move for $i = 0$, are not in any way restricted in the r -local Hintikka game, our minimality setup guarantees that ψ starts with a quantifier, and so $q + 1 \geq 1$ and $i \geq 1$. If this leading quantifier of ψ was \bar{Q} , then we would again get a contradiction to the minimality of q . Therefore,

$$(1) \quad \psi(p_1, \dots, p_i) \equiv Qx_{i+1}. \psi'(p_1, \dots, p_i, x_{i+1}).$$

Note that q is the quantifier rank of ψ' . If $q = 0$, then ψ' is actually quantifier-free and Player Q can make his move x_{i+1} with the element $p' \in R_{r_0}^{D_0}(p_1, \dots, p_i)$, $r_0 = 2$, as in Lemma 3.5. Since

this is a contradiction to Player Q losing $\mathcal{G}_r(\mathcal{P}, \psi)$ from $\psi(p_1, \dots, p_i)$, we may further assume that $q \geq 1$.

Let $p \in P$ be such that Player Q wins the game $\mathcal{G}(\mathcal{P}, \varphi)$ from (1) $\psi(p_1, \dots, p_i)$ by moving to the position $\psi'(p_1, \dots, p_i, p)$, and assume that p is chosen maximal with respect to \leq^P with this property. By our minimal choice of q , the statement of Lemma 4.2 holds for the game $\mathcal{G}(\mathcal{P}, \varphi)$ from the position $\psi'(p_1, \dots, p_i, p)$, and so

- (2) Player Q wins the r -local game $\mathcal{G}_r(\mathcal{P}, \psi')$ from the initial position $\psi'(p_1, \dots, p_i, p)$.

Consequently,

- (3) $p \notin R_{r_q - r_{q-1}}^{D_q}(p_1, \dots, p_i)$

since, otherwise, Player Q would win also the r -local game $\mathcal{G}_r(\mathcal{P}, \psi)$ from $\psi(p_1, \dots, p_i)$ which is not the case by our assumption.

Let $p' \in C(p)$ be the bottommost element such that $p \leq^P p'$ and $\tau_q(p') = \tau_q(p)$. Observe that such p' does exist since, otherwise, D_q would contain an arc from the topmost element of $C(p)$ to p with label ‘down’, and hence $p \in R_2^{D_q}(p_1)$ contradicting (3). By our maximal choice of p with respect to \leq^P , Player Q loses the game $\mathcal{G}(\mathcal{P}, \varphi)$ from $\psi(p_1, \dots, p_i)$ after moving to the position $\psi'(p_1, \dots, p_i, p')$. In other words, Player \bar{Q} wins the game $\mathcal{G}(\mathcal{P}, \varphi)$ from the position $\psi'(p_1, \dots, p_i, p')$ and, by our choice of a counterexample with minimum q ;

- (4) Player \bar{Q} wins the r -local game $\mathcal{G}_r(\mathcal{P}, \psi')$ from the initial position $\psi'(p_1, \dots, p_i, p')$.

Now, we employ Corollary 3.7 for $s = q - 1$, $k = i$ and $S = R_{r_{q-1}}^{D_{q-1}}(p_1, \dots, p_i)$. By (3), we hence get that there exists a color-preserving isomorphism map f between the induced colored subposets $\mathcal{P}[S \cup P_{q-1}(p)]$ and $\mathcal{P}[S \cup P_{q-1}(p')]$. By Lemma 4.1, all the elements e played in any play of an r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \psi')$ from an initial position $\psi'(p_1, \dots, p_i, p')$ belong to the set $R_{r_{q-1}}^{D_{q-1}}(p_1, \dots, p_i, p') = S \cup P_{q-1}(p')$.

The latter finding implies that every play of the r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \psi')$ from $\psi'(p_1, \dots, p_i, p')$ can be duplicated, with the same outcome, in the same game from $\psi'(p_1, \dots, p_i, p)$ via the isomorphism map f^{-1} . Therefore, by (4), Player \bar{Q} wins the r -local game $\mathcal{G}_r(\mathcal{P}, \psi')$ from the initial position $\psi'(p_1, \dots, p_i, p)$, too. However, this contradicts (2), and so there cannot be a counterexample to the statement of the lemma. \square

Remark 4.3. Notice that it is actually not necessary to explicitly use Lemma 3.5 in the proof of Lemma 4.2—with a slightly modified setting this base case comes out “for free”. Though, we think that the current proof with Lemma 3.5 is easier to read and to understand.

We can now easily formulate and prove the main result:

Theorem 4.4. *Let $\mathcal{P} = (P, \leq^P)$ be a poset, associated with $\lambda : P \rightarrow \Lambda$ where Λ is a finite set of colors, and let φ be an FO sentence in negation normal form. The existential player has a winning strategy in the r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \varphi)$ if, and only if, $\mathcal{P} \models \varphi$.*

Proof. By Proposition 2.2, it is enough to prove the following:

- Player Q , where $Q \in \{\exists, \forall\}$, has a winning strategy in the Hintikka game $\mathcal{G}(\mathcal{P}, \varphi)$ if, and only if, Player Q has a winning strategy in the r -local Hintikka game $\mathcal{G}_r(\mathcal{P}, \varphi)$.

The “only if” direction follows from Lemma 4.2 for $i = 0$, and so we deal with the “if” direction. Assume that we have a counterexample minimizing the quantifier rank of φ , i.e., that Player Q wins $\mathcal{G}_r(\mathcal{P}, \varphi)$ but does not win $\mathcal{G}(\mathcal{P}, \varphi)$.

If the only quantifiers in φ are Q , then the r -local winning strategy of Player Q can be used “as is” in the game $\mathcal{G}(\mathcal{P}, \varphi)$. Hence, denoting by \bar{Q} the opposite quantifier to Q , the sentence φ contains a subformula $\psi \equiv \bar{Q}y. \psi'(y)$, such that all occurrences of the quantifier \bar{Q} in φ are contained within ψ . Let $\psi(p_1, \dots, p_i)$ be a position reached by any winning strategy of Player Q in the r -local game $\mathcal{G}_r(\mathcal{P}, \varphi)$. Since Player Q loses $\mathcal{G}(\mathcal{P}, \varphi)$ in general, Player \bar{Q} wins the ordinary game $\mathcal{G}(\mathcal{P}, \varphi)$ from the position $\psi(p_1, \dots, p_i)$.

By Lemma 4.2 again, Player \bar{Q} hence wins the r -local game $\mathcal{G}_r(\mathcal{P}, \psi)$ from the initial position $\psi(p_1, \dots, p_i)$. However, this contradicts the assumption that the strategy of Player Q in $\mathcal{G}_r(\mathcal{P}, \varphi)$ containing the position $\psi(p_1, \dots, p_i)$ is winning for Player Q . \square

The following is then a straightforward observation.

Corollary 4.5. *For a poset \mathcal{P} let $\mathcal{T}_i(\mathcal{P})$ denote the set of types τ_i of rank i occurring in \mathcal{P} , i.e. $\mathcal{T}_i(\mathcal{P}) := \{\tau_i(p) : p \in P\}$, where τ_i is as in Definition 3.1. Assume $\mathcal{P}^1, \mathcal{P}^2$ are colored posets and φ an FO sentence of quantifier rank q in negation normal form. If $\mathcal{T}_{q-1}(\mathcal{P}^1) = \mathcal{T}_{q-1}(\mathcal{P}^2)$, then it holds; $\mathcal{P}^1 \models \varphi$ if and only if $\mathcal{P}^2 \models \varphi$.*

Proof. Let the types of elements of \mathcal{P}^1 and \mathcal{P}^2 be denoted by τ_i^1 and τ_i^2 , respectively, and the associated digraphs by D_i^1 and D_i^2 .

From Theorem 4.4, it is enough to prove that the existential player wins the r -local game $\mathcal{G}_r(\mathcal{P}^1, \varphi)$ if, and only if, he wins $\mathcal{G}_r(\mathcal{P}^2, \varphi)$. For that we have to analyze the game up to the first quantifier move, and since the starting conjunction or disjunction moves are irrelevant for this analysis, we may without loss of generality assume that $\varphi \equiv Qy. \psi(y)$ where $Q \in \{\exists, \forall\}$.

If the claim (that the existential player wins $\mathcal{G}_r(\mathcal{P}^1, \varphi)$ iff he wins $\mathcal{G}_r(\mathcal{P}^2, \varphi)$) is not true, then Player Q wins precisely one of the two considered games, say, the r -local game $\mathcal{G}_r(\mathcal{P}^1, \varphi)$. Let a winning move of Player Q in $\mathcal{G}_r(\mathcal{P}^1, \varphi)$ be the position $\psi(p_1)$ where $p_1 \in P^1$. By the assumption $\mathcal{T}_{q-1}(\mathcal{P}^1) = \mathcal{T}_{q-1}(\mathcal{P}^2)$, there hence exists $p_2 \in P^2$ such that $\tau_{q-1}^1(p_1) = \tau_{q-1}^2(p_2)$. Let the latter be witnessed by an isomorphism $\iota : P_{q-2}^1(p_1) \rightarrow P_{q-2}^2(p_2)$ between the structures $\mathcal{A}_{q-2}^{P^1}(p_1)$ and $\mathcal{A}_{q-2}^{P^2}(p_2)$ which is, in particular, a poset isomorphism between the induced colored subposets $\mathcal{P}^1[P_{q-2}^1(p_1)]$ and $\mathcal{P}^2[P_{q-2}^2(p_2)]$.

By Lemma 4.1, for the formula ψ of quantifier rank $q - 1$ and for $j = 1$, every reachable position in the r -local game $\mathcal{G}_r(\mathcal{P}^1, \varphi)$ belongs to $P_{q-2}^1(p_1)$. Consequently, Player Q could win also the r -local game $\mathcal{G}_r(\mathcal{P}^2, \varphi)$, by copying his winning strategy under the isomorphism map ι . This contradiction finishes the proof. \square

Theorem 4.6. *Let $\mathcal{P} = (P, \leq^{\mathcal{P}})$ be a poset of width w , with elements colored by $\lambda : P \rightarrow \Lambda$ where Λ is a finite set, and let φ be an FO sentence in negation normal form. There is an algorithm which decides whether $\mathcal{P} \models \varphi$ in FPT time $f(w, \varphi) \cdot \|\mathcal{P}\|^2$.*

Proof. According to Corollary 4.5, it is enough to know the set $\mathcal{T}_{q-1}(\mathcal{P})$ in order to decide whether $\mathcal{P} \models \varphi$. We thus proceed the algorithm in two steps:

1. We compute the set $\mathcal{T}_{q-1}(\mathcal{P})$ (of rank- $(q - 1)$ types).
2. We decide whether $\mathcal{P} \models \varphi$ using the set from Step 1.

First of all, we show that the set of all possible types of a given rank is finite if the poset width w is bounded. Since each type is a (sub)digraph of bounded radius, it is enough to argue that the out-degrees in D_i are bounded. Indeed, by induction, the outdegree in D_0 is $4w|\Lambda|$. In

D_{i+1} , the outdegree is bounded from above by a function of w and the number of all possible types of rank $i + 1$, which is finite by the inductive assumption for D_i .

Therefore, Step 2 is a finite problem and we may decide whether $\mathcal{P} \models \varphi$ by a brute-force evaluation of φ on each member of $\mathcal{T}_{q-1}(\mathcal{P})$. This takes time $f'(w, \varphi)$. As for Step 1, we start with computing a chain partition of width w , in time $g(w) \cdot \|\mathcal{P}\|^2$ by Proposition 2.1. We then proceed exactly along the iterations of constructive Definition 3.1. Since the number of possible types is finite, this computation takes time at most $f''(w, \varphi) \cdot \|\mathcal{P}\|^2$; by traversing, in every iteration D_i for $i = 0, 1, \dots, q - 2$, for each $p \in P$ every chain of \mathcal{P} and finding the appropriate out-neighbors. \square

5 Application to Interval Graphs

Besides the very successful story of FPT FO model checking on sparse graph classes, culminating with the ultimate and outstanding result of Grohe, Kreutzer, and Siebertz [15], only a few such results have been published for dense graph classes (especially, for graph classes which cannot be easily interpreted in nowhere dense classes). One of such notable papers is [13], dealing with FO model checking on interval graphs.

It has been shown [13] that FO model checking of interval graphs is FPT when the intervals are restricted to have lengths from a fixed finite set of reals (Corollary 5.2), while the problem is W-hard whenever the intervals are allowed to have lengths from any dense subset of a positive-length interval of reals. We will demonstrate the strength and usefulness of our main result by giving a rather short and straightforward derivation of the FPT result of [13] from our Theorem 4.6.

A graph G is an *interval graph* if there exists a set \mathcal{I} of intervals on the real line such that $V(G) = \mathcal{I}$ and $E(G)$ is formed by the intersecting pairs of intervals. For a set L of reals, a set \mathcal{I} of intervals is called an *L -interval representation* if every interval from \mathcal{I} has its length in L . This notion generalizes well-studied *unit interval graphs* (where all interval lengths are 1), which are also known under the name of proper interval graphs. A set \mathcal{I} of intervals is called a *proper interval representation* if there is no pair of intervals $J_1, J_2 \in \mathcal{I}$ such that J_1 is strictly contained in J_2 ($J_1 \subsetneq J_2$). We call \mathcal{I} a *k -fold proper interval representation* if there exists a partition $\mathcal{I} = \mathcal{I}_1 \cup \dots \cup \mathcal{I}_k$ such that each \mathcal{I}_j is a proper interval representation for $j = 1, \dots, k$.

Theorem 5.1. *Let φ be a graph FO sentence. Assume G is an interval graph given along with its k -fold proper interval representation \mathcal{I} . Then the FO model checking problem $G \models \varphi$, parameterized by k and φ , is FPT.*

Proof. First of all, we prove that we can, without loss of generality, assume that no two ends of intervals from \mathcal{I} coincide. This can be achieved by a tiny perturbation of the intervals, as follows. Let the given interval representation of G be $\mathcal{I} = \{[a_i, b_i] : i = 1, \dots, n\}$, where $n = |V(G)|$. Let $D = \bigcup_{i=1}^n \{a_i, b_i\}$ be the set of all interval ends. Choose $\varepsilon > 0$ such that the least positive difference in the set D is greater than 2ε , and define

$$\mathcal{I}' := \left\{ \left[a_i + \varepsilon \frac{i}{n}, b_i + \varepsilon \left(1 + \frac{i}{n} \right) \right] : i = 1, \dots, n \right\}.$$

Then, clearly, the intersection graph of \mathcal{I}' is isomorphic to G , no two interval ends in \mathcal{I}' coincide, and two intervals of \mathcal{I}' are in a proper inclusion only if the same is true already in \mathcal{I} .

Let $\mathcal{I} = \mathcal{I}_1 \cup \dots \cup \mathcal{I}_k$ be such that each \mathcal{I}_j is a proper interval representation for $j = 1, \dots, k$. Let $P := D \cup \mathcal{I}$. We define a poset $\mathcal{P} = (P, \leq^{\mathcal{P}})$ as follows:

- for numbers $d_1, d_2 \in D$ it is $d_1 \leq^{\mathcal{P}} d_2$ iff $d_1 \leq d_2$,

- for $j \in \{1, \dots, k\}$ and intervals $J_1, J_2 \in \mathcal{I}_j$, it is $J_1 \leq^{\mathcal{P}} J_2$ iff J_1 is not to the right of J_2 , and
- for every interval $J = [a, b] \in \mathcal{I}$ and every $d \in D$, it is $J \leq^{\mathcal{P}} d$ iff $d \geq b$, and $d \leq^{\mathcal{P}} J$ iff $d \leq a$.

The set \mathcal{P} can be partitioned into $k + 1$ chains; D and $\mathcal{I}_1, \dots, \mathcal{I}_k$. Hence the width of \mathcal{P} is at most $k + 1$.

Finally, we show how the poset \mathcal{P} can be used to decide whether $G \models \varphi$, applying a so called FO interpretation of G in \mathcal{P} . Specifically, we construct a poset FO sentence ψ , depending only on φ , such that $\mathcal{P} \models \psi$ if and only if $G \models \varphi$. The rest then follows from Theorem 4.6.

Let $D(x)$ be a unary predicate (a color) on \mathcal{P} identifying the elements of $D \subseteq \mathcal{P}$. Then, the proposition $\neg D(x)$ simply identifies the domain (vertex set) of G within \mathcal{P} in our interpretation. The crucial part of our construction of ψ is to express by an FO formula $\beta(J_1, J_2)$ that arbitrary intervals $J_1, J_2 \in \mathcal{I}$ intersect each other (i.e., form an edge of G). Note that if intervals J_1, J_2 do not intersect, then J_2 is to the right of the right end of J_1 or vice versa. Then, by our definition of \mathcal{P} , J_2 is to the right of the right end of J_1 if, and only if, there exists $d \in D$ such that $J_1 \leq^{\mathcal{P}} d \leq^{\mathcal{P}} J_2$ (d lies “between” J_1 and J_2). Hence we can express the edge relation of G as

$$\beta(x, y) \equiv \forall d [D(d) \rightarrow ((\neg x \leq^{\mathcal{P}} d \vee \neg d \leq^{\mathcal{P}} y) \wedge (\neg y \leq^{\mathcal{P}} d \vee \neg d \leq^{\mathcal{P}} x))],$$

meaning that there exists no $d \in D$ such that d lies “between” the intervals x and y (or, y and x) of \mathcal{I} .

The construction of ψ then proceeds by structural induction on the structure of φ as follows. Each occurrence of atomic $edge(x, y)$ is replaced with $\beta(x, y)$. Other atomic propositions and logic connectives are simply copied from φ to ψ . Each quantifier $\exists x \sigma(x)$ is replaced with $\exists x (\neg D(x) \wedge \sigma(x))$, and $\forall x \sigma(x)$ with $\forall x (\neg D(x) \rightarrow \sigma(x))$. All this construction is achieved in polynomial time. Based on previous arguments, it is routine to verify that $\mathcal{P} \models \psi \iff G \models \varphi$. \square

We can now match the main result of [13], except the precise runtime:

Corollary 5.2 (Ganian et al. [13]). *For every finite set L of reals, the FO model checking problem of L -interval graphs (given alongside with an L -interval representation) is FPT when parameterized by the FO sentence φ and $|L|$.*

Proof. Let \mathcal{I} be a given L -interval representation, and set $k := |L|$. We partition $\mathcal{I} = \mathcal{I}_1 \cup \dots \cup \mathcal{I}_k$ such that each \mathcal{I}_i contains intervals of the same length (from L), and then apply Theorem 5.1. \square

Although, the formulation of Theorem 5.1 is more general than [13]. We can, for instance, in the same way derive fixed-parameter tractability also for FO model checking of well-studied *k-proper interval graphs*, introduced in [20] as those having an interval representation such that no interval is properly contained in more than k other intervals.

6 Conclusions

Our result can be seen as an initial step towards an understanding of the complexity of FO model checking on non-sparse classes of structures and we hope that the techniques developed here will be useful for future research in this direction, e.g., to investigate the complexity of FO model checking on other algebraic structures such as finite groups and lattices as suggested by Grohe [15], as well as on other dense graph classes, to which the established “locality” tools of finite model theory do not apply.

The result may also be used directly towards establishing fixed-parameter tractability for FO model checking of other graph classes. Given the ease with which it implies the otherwise non-trivial result on interval graphs [13], it is a natural to ask which other (dense) graph classes can be interpreted in posets of bounded width.

References

- [1] S. Bova, R. Ganian, and S. Szeider. Model checking existential logic on partially ordered sets. In *CSL-LICS'14*. ACM, 2014. Article No. 21.
- [2] S. Bova, R. Ganian, and S. Szeider. Quantified conjunctive queries on partially ordered sets. In *IPEC'14*, volume 8894 of *LNCS*, pages 122–134. Springer, 2014.
- [3] B. Courcelle. The monadic second-order logic of graphs I: Recognizable sets of finite graphs. *Inform. and Comput.*, 85:12–75, 1990.
- [4] A. Dawar, M. Grohe, and S. Kreutzer. Locally excluding a minor. In *LICS'07*, pages 270–279. IEEE Computer Society, 2007.
- [5] R. Downey and M. Fellows. *Parameterized complexity*. Monographs in Computer Science. Springer, 1999.
- [6] Z. Dvořák, D. Král', and R. Thomas. Deciding first-order properties for sparse graphs. In *FOCS'10*, pages 133–142. IEEE Computer Society, 2010.
- [7] E. Grädel et al. *Finite Model Theory and Its Applications (Texts in Theoretical Computer Science. An EATCS Series)*. Springer, 2005.
- [8] S. Felsner, V. Raghavan, and J. Spinrad. Recognition algorithms for orders of small width and graphs of small dilworth number. *Order*, 20(4):351–364, 2003.
- [9] J. Flum and M. Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.*, 31(1):113–145, 2001.
- [10] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [11] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001.
- [12] J. Gajarský, P. Hliněný, J. Obdržálek, and S. Ordyniak. Faster existential FO model checking on posets. In *ISAAC'14*, volume 8889 of *LNCS*, pages 441–451. Springer, 2014.
- [13] R. Ganian, P. Hliněný, D. Král', J. Obdržálek, J. Schwartz, and J. Teska. FO model checking of interval graphs. In *ICALP 2013, Part II*, volume 7966 of *LNCS*, pages 250–262. Springer, 2013.
- [14] M. Grohe. Logic, graphs, and algorithms. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(091), 2007.
- [15] M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. In *STOC'14*, pages 89–98. ACM, 2014.
- [16] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [17] S. Kreutzer. Algorithmic meta-theorems. In *Finite and Algorithmic Model Theory*, volume 379 of *London Math. Soc. Lecture Note Ser.*, chapter 5, pages 177–270. Oxford University Press, 2011.
- [18] D. Lokshantov, D. Marx, and S. Saurabh. Lower bounds based on the exponential time hypothesis. *Bulletin of the EATCS*, 105:41–72, 2011.
- [19] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*, volume 31 of *Oxford Lecture Ser. Math. Appl.* OUP, 2006.
- [20] A. Proskurowski and J. A. Telle. Classes of graphs with restricted interval models. *Discrete Math. Theor. Comput. Sci.*, 3(4):167–176, 1999.
- [21] D. Seese. Linear time computable problems and first-order descriptions. *Math. Structures Comput. Sci.*, 6(6):505–526, 1996.