

Model Checking Lower Bounds for Simple Graphs

Michael Lampis*

KTH Royal Institute of Technology
mlampis@kth.se

Abstract. A well-known result by Frick and Grohe shows that deciding FO logic on trees involves a parameter dependence that is a tower of exponentials. Though this lower bound is tight for Courcelle’s theorem, it has been evaded by a series of recent meta-theorems for other graph classes. Here we provide some additional non-elementary lower bound results, which are in some senses stronger. Our goal is to explain common traits in these recent meta-theorems and identify barriers to further progress.

More specifically, first, we show that on the class of threshold graphs, and therefore also on any union and complement-closed class, there is no model-checking algorithm with elementary parameter dependence even for FO logic. Second, we show that there is no model-checking algorithm with elementary parameter dependence for MSO logic even restricted to paths (or equivalently to unary strings), unless $\text{EXP}=\text{NEXP}$. As a corollary, we resolve an open problem on the complexity of MSO model-checking on graphs of bounded max-leaf number. Finally, we look at MSO on the class of colored trees of depth d . We show that, assuming the ETH, for every fixed $d \geq 1$ at least $d+1$ levels of exponentiation are necessary for this problem, thus showing that the $(d+1)$ -fold exponential algorithm recently given by Gajarský and Hliněný is essentially optimal.

1 Introduction

Algorithmic meta-theorems are general statements establishing tractability for a whole class of problems (often defined by expressibility in a certain logic) in some class of inputs (usually a family of graphs). By far the most famous and celebrated theorem in this area is a twenty-year old result due to Courcelle [2] which states that all problems expressible in monadic second-order logic (MSO_2) are linear-time solvable on graphs of bounded treewidth. Much work has been devoted in recent years to proving stronger and stronger meta-theorems in this spirit, often extending Courcelle’s theorem to other graph classes (see e.g. [3, 7, 5] or [12, 13] for some great surveys).

The most often cited drawback of Courcelle’s theorem has to do with the “hidden constant” in the algorithm’s linear running time. It is clear

* Research supported by ERC Grant 226203

that the running time must somehow depend on the input formula and the graph's treewidth, but the dependence given in Courcelle's theorem is in the worst case a tower of exponentials whose height grows with the size of the formula. Unfortunately, this cannot be avoided: Frick and Grohe [8] proved that the parameter dependence has to be non-elementary even if one restricts the problem severely by just looking at properties expressible in first-order logic on trees.

This lower bound result, though quite devastating, has proven very fruitful and influential: several papers have appeared recently with the explicit aim of proving meta-theorems which evade it, and thus achieve a much better dependence on the parameters. Specifically, in [15] algorithmic meta-theorems with an elementary parameter dependence are shown for vertex cover, max-leaf number and the newly defined neighborhood diversity. A meta-theorem for twin cover was shown by Ganian [10]. In addition, meta-theorems were shown for tree-depth by Gajarský and Hliněný [9] and for the newly defined shrub-depth (which generalizes neighborhood diversity and twin cover) by Ganian et al. [11].

Thus, together with improved meta-theorems, these papers give a new crop of graph complexity measures, some more general than others. It becomes a natural question how much progress we can hope to achieve this way, that is, how far this process of defining more and more general "graph widths" can go on before hitting some other natural barrier that precludes an elementary parameter dependence. Is simply avoiding the class of all trees enough?

This is exactly the question we try to answer in this paper. Towards this end we try to give hardness results for graph families which are as simple as possible. Perhaps most striking among them is a result showing that not only is avoiding all trees not enough but in fact it is necessary to avoid the much smaller class of uncolored paths if one hopes for an elementary parameter dependence. As an example application, this almost immediately rules out the existence of meta-theorems with elementary parameter dependence in any induced-subgraph-closed graph class with unbounded diameter and any edge-subdivision-closed graph class. This explains why all recently shown meta-theorems we mentioned work on classes which are closed under induced subgraphs but have bounded diameter and are not closed under edge subdivisions.

Our results can be summarized as follows. First, a non-elementary lower bound for model checking FO logic on threshold graphs is shown. In a sense, this is a natural analogue of the lower bound for trees to the realm of clique-width, since threshold graphs are known to have the smallest

possible clique-width. The proof is relatively simple and consists mostly of translating a similar lower bound given in [8] for FO model checking on binary words. However, the main interest of this result is that as a corollary we show that the complexity of FO model checking is non-elementary for any graph class closed under disjoint union and complement. This explains why, though some of the recent meta-theorems work on complement-closed graph classes (e.g. neighborhood diversity, shrub-depth) and some work on union-closed graph classes (e.g. tree-depth), no such meta-theorem has been shown for a class that has both properties.

Our second result is that model checking MSO logic on uncolored paths (or equivalently on unary strings) has a non-elementary parameter dependence. This is the most technically demanding of the results of this paper, and it is proved under the assumption $\text{EXP} \neq \text{NEXP}$. The proof consists of simulating the workings of a non-deterministic Turing machine via an MSO question on a path. Though the idea of simulating Turing machines has appeared before in similar contexts [14], because the graphs we have here are very restricted we face a number of significant new challenges. The main tool we use to overcome them, which may be of independent interest, is an MSO formula construction that compares the sizes of ordered sets while using an extremely small number of quantifiers. In the end, this result strengthens both non-elementary MSO lower bounds given in [8] (for trees and for binary strings), modulo a slightly stronger complexity assumption. It also resolves the complexity of MSO model checking for max-leaf number, which was left open in [15]. As an added corollary, we give an alternative, self-contained proof of a result from [3], stating that MSO_2 model checking is not in XP for cliques unless $\text{EXP} = \text{NEXP}$.

Finally, we study one of the recent positive results in this area by considering the problem of model-checking MSO logic on rooted colored trees of height d . This is an especially interesting problem, since the $(d + 1)$ -fold exponential algorithm of [9] is the main tool used in the meta-theorems of both [9] and [11]. We show that, assuming the ETH, any algorithm needs at least $d + 1$ levels of exponentiation, and therefore the algorithm of [9] is essentially optimal.

2 Preliminaries

The basic problem we are concerned with is model-checking: We are given a formula ϕ (in some logic) and a structure S (usually a graph or a string) and must decide if $S \models \phi$, that is, if S satisfies the property described by ϕ .

Due to space constraints, we do not give a full definition of FO and MSO logic here (see e.g. [8]). Let us just briefly describe some conventions. We use lower-case letters to denote singleton (FO) variables, and capitals to denote set variables. When the input is a graph, we assume the existence of an $E(x, y)$ predicate encoding edges; when it's a string a \prec predicate encodes a total ordering; when it's a rooted tree a $C(x, y)$ predicate encodes that x is a child of y . Sometimes the input also has a set of colors (also called labels). For each color c we are given a unary predicate $P_c(x)$. When the input is an uncolored graph that consists of a single path it is possible to simulate the \prec predicate by picking one endpoint of the path arbitrarily (call it s) and saying that $x \prec y$ if all paths from s to y contain x . Thus, model-checking MSO logic on uncolored paths is at least as hard as it is on unary strings. In most of the paper when we talk about MSO logic for graphs we mean MSO_1 , that is, with quantification over vertex sets only. An exception is Corollary 4 which talks about MSO_2 logic, which allows edge set quantifiers.

A graph is a threshold graph ([1]) if it can be constructed from K_1 by repeatedly adding *union* vertices (not connected to any previous vertex) and *join* vertices (connected to all previous vertices), one at a time. Thus, a threshold graph can be described by a string over the alphabet $\{u, j\}$. A graph is a cograph if it is K_1 , or it is a disjoint union of cographs, or it is the complement of a cograph. It is not hard to see that threshold graphs are cographs. From the definition it follows that any class of graphs that contains K_1 and is closed under disjoint union and complement contains all cographs; if it is closed under the union and join operations it contains all threshold graphs.

All logarithms are base two. We define $\exp^{(k)}(n)$ as follows: $\exp^{(0)}(n) = n$ and $\exp^{(k+1)}(n) = 2^{\exp^{(k)}(n)}$. Then $\log^{(k)} n$ is the inverse of $\exp^{(k)}(n)$. Finally, $\log^* n$ is the minimum i such that $\log^{(i)} n \leq 1$.

3 Threshold Graphs

As mentioned, Frick and Grohe [8] showed that there is no FPT model-checking algorithm for FO logic on trees with an elementary dependence on the formula size, under standard complexity assumptions. In many senses this is a great lower bound result, because it matches the tower of exponentials that appears in the running time of Courcelle's theorem, while looking both at a much simpler logic (FO rather than MSO_2) and at the class of graphs with the smallest possible treewidth, namely trees.

Courcelle, Makowsky and Rotics [3] have given an extension of Courcelle’s theorem to MSO_1 logic for clique-width. The parameter dependence is again a tower of exponentials and, since trees have cliquewidth at most 3 ([4]), we already know that this cannot be avoided even for graphs of constant clique-width. Here we will slightly strengthen this result, showing that the non-elementary dependence cannot be avoided even on cographs, the class of graphs that has the smallest possible clique-width (that is, clique-width 2) without being trivial. We will heavily rely on a lower bound, due again to Frick and Grohe, on the complexity of model checking on binary strings.

One interesting consequence of the lower bound we give for cographs is that it precludes the existence of an FPT algorithm with elementary parameter dependence for any graph class that satisfies two simple properties: closure under disjoint unions and closure under complement. The reason for this is that if a class is closed under both of these operations and it contains the single-vertex graph, then it must contain all cographs (we will also show that the assumption that K_1 is in the class is not needed). This observation helps to explain why, though some of the recent elementary model-checking algorithms which have appeared work on union-closed graph classes, and some work on complement-closed graph classes, no such algorithms are known for classes with both properties.

The proof we present here is relatively simple and it relies on the following theorem.

Theorem 1 ([8]).

Unless $\text{FPT}=\text{AW}[]$, for any constant c and any elementary function f there is no model-checking algorithm for FO logic on binary words which given a formula ϕ and a word w decides if $w \models \phi$ in time at most $f(\phi)|w|^c$.*

We will reduce this problem to FO model checking on threshold graphs. This is quite natural, since the definition of threshold graphs gives a straightforward correspondence between graphs and strings.

Theorem 2. *Unless $\text{FPT}=\text{AW}[*]$, for any constant c and any elementary function f there is no model-checking algorithm for FO logic on connected threshold graphs which given a formula ϕ and such a graph G decides if $G \models \phi$ in time at most $f(\phi)|G|^c$.*

Proof (sketch).

We can encode an arbitrary binary string with a threshold graph by encoding each 0 by two vertices uj and each 1 by three vertices ujj . We also need to encode a FO formula for strings to one for graphs. Existential

quantification in the string is simulated by existential quantification in the graph where we ask that the vertex selected is a union vertex (this can be expressed in FO logic). The predicates $x \prec y$ and $P_1(x)$ can be implemented by checking if there are join vertices connected to x but not y in the first case, and checking for two join vertices connected to x but not later union vertices in the second. \square

Corollary 1. *Let \mathcal{C} be a graph class that is closed under disjoint union and complement, or under disjoint union and join. Unless $FPT=AW[*]$, for any constant c and any elementary function f there is no model-checking algorithm for FO logic on \mathcal{C} which given a formula ϕ and a graph $G \in \mathcal{C}$ decides if $G \models \phi$ in time at most $f(\phi)|G|^c$.*

4 Paths, Unary Strings

The main result of this section is a reduction proving that, under the assumption that $EXP \neq NEXP$, there is no FPT model-checking algorithm for MSO logic with an elementary parameter dependence on the formula even on graphs that consist of a single path, or equivalently, on unary strings. As a consequence, this settles the complexity of MSO model-checking on graphs with bounded max-leaf number, a problem left open in [15], since paths have the smallest possible max-leaf number. Until now a similar result was known only for the much richer class of binary strings (or equivalently colored paths), under the weaker assumption that $P \neq NP$ [8]. It is somewhat surprising that we are able to extend this result to uncolored paths, because in this case the size of the input is exponentially blown-up compared to a reasonable encoding. One would expect this to make the problem easier, but in fact, it only makes it more complicated to establish hardness.

Indeed, one of the main hurdles in proving a lower bound for MSO on unary strings, or paths, is information-theoretic. Normally, one would start with an NP-hard problem, and reduce to a model-checking instance with a very small formula ϕ . But, because the path we construct can naturally be stored with a number of bits that is logarithmic in its size, in order to encode n bits of information from the original instance into the new instance we need to construct a path of exponential size. Thus, a polynomial-time reduction seems unlikely and this is the reason we end up using the assumption that $EXP \neq NEXP$, instead of $P \neq NP$.

Our approach is to start from the prototypical NEXP-complete problem: given n bits of input for a non-deterministic Turing machine that runs in time 2^{n^k} , does the machine accept? We will use the input path to

simulate the machine's tape and then ask for a subset of the vertices of this path that corresponds to cells in the tape where 1 is written. Thus, what we need at this point is an MSO formula that checks if the chosen vertices encode a correct accepting computation.

Of course, to describe a machine's computation in MSO logic a significant amount of machinery will be needed. We note that, though the approach of simulating a Turing machine with an MSO formula has been used before (e.g. [14]), the problem here is significantly more challenging for two reasons: first, unlike previous cases the input here is uncolored, so it is harder to encode arbitrary bits; and second, there are (obviously) no grid-like minors in our graph, so it's harder to encode the evolution of a machine's tape, and in particular to identify vertices that correspond to the same tape cell in different points in time.

Our main building block to overcome these problems is an MSO construction which compares the sizes of paths (or generally, ordered sets) of size n with very few (roughly $2^{O(\log^* n)}$) quantifiers. This construction may be of independent interest in the context of the counting power of MSO logic. We first describe how to build this formula, then use it to obtain other basic arithmetic operations (such as exponentiation and division) and finally explain how they all fit together to give the promised result.

4.1 Measuring Long Paths with Few Quantifiers

To keep the presentation simple we will concentrate on the model-checking problem on unary strings; formulas for MSO on paths can easily be constructed as explained in section 2. We therefore assume that there is a predicate \prec which gives a total ordering of all elements.

Let us now develop our basic tool, which will be an MSO formula $eq_L(P_1, P_2)$, where P_1, P_2 are free set variables. The desired behavior of the formula is that if $|P_1| = |P_2|$ and $|P_1| \leq L$ then the formula will be true, while on the other hand whenever the formula is true it must be the case that $|P_1| = |P_2|$. In other words, the formula will always correctly identify equal sets with size up to L , and it will never identify two unequal sets as equal (it may however be false for two equal sets larger than L). Our main objective is to achieve this with as few quantifiers as possible.

We will work inductively. It should be clear that for very small values of L (say $L = 4$) it is possible to compare sets of elements with size at most L with a constant number of set and vertex quantifiers and we can simply make the formula false if one set has more than 4 elements. So, suppose that we have a way to construct the desired formula for some L . We will show how to use it to make the formula $eq_{L'}$, where $L' \geq L \cdot 2^L$.

If our recursive definition of eq_L uses a constant number of copies of eq_L then in the end we will have $|eq_L| = 2^{O(\log^* L)}$, because for each level of exponentiation we blow up the size of a formula by a constant factor. This will be sufficiently small to rule out a non-elementary parameter dependence.

Let us now give a high-level description of the idea, by concentrating first on the set P_1 . We will select a subset of P_1 , call it Q_1 , and this naturally divides P_1 into sections, which are defined as maximal sets of vertices of P_1 , consecutive in the ordering, with the property that either all or none of their vertices belong in Q_1 . We will make sure that all sections have length L , except perhaps the last, which we call the remainder (see Figure 1). It is not hard to see that this structure can be imposed if the predicate eq_L is available. We do the same for P_2 and now we need to verify that the two remainders have the same length (easy with eq_L) and that we have the same number of sections on P_1 and P_2 .



Fig. 1. An example of the counting structure imposed on a set for $L = 4$. Assume that the elements of the set are displayed here in the ordering given by \prec . We select a set of elements (indicated by boxes) to create sections of elements of size L , and a remainder set (indicated by a dashed box). We then select some elements from each section appropriately (indicated by solid grey filling) so that we simulate counting in binary. For $L = 4$ this method can count up to length $L2^L = 64$ before overflowing.

We need to count the number of sections this structure creates on P_1 (which may be up to 2^L). We select another subset of P_1 , call it B_1 . The intuition here is that selecting B_1 corresponds to writing a binary number on each section, by interpreting positions selected in B_1 as 1 and the rest as 0. We will now need to make sure that each section encodes the binary number that is one larger than the number encoded by the immediately preceding section. This is achievable by using eq_L to locate the elements that represent the same bit positions. We also make sure that there was no overflow in the counting and that counting started from zero, that is, all sections have some vertex not in B_1 and the first has no vertices in B_1 .

Finally, assuming that the above counting structure is correctly imposed on both P_1 and P_2 all that is left is to take the last sections on both and compare them. If the same binary number is encoded on both then $|P_1| = |P_2|$.

A formal definition for eq_L and all other formulas of this section is given in the appendix.

Lemma 1. *Let $L > 2$ be a power of two. Then we can define a formula $eq_L(P_1, P_2)$ such that the formula is true if and only if $|P_1| = |P_2| < L \log L$. Furthermore $|eq_L| = 2^{O(\log^* L)}$.*

Before we go on, we will also need formulas to perform some slightly more complicated arithmetic operations than simply counting. In particular, we will need a formula $exp_L(P_1, P_2)$, which will be true if $|P_2| = 2^{|P_1|}$. The trick we use for this is shown in Figure 2.



Fig. 2. An example where the set on the left has size 4 and we verify that the set on the right has size 2^4 . First we select a subset of elements on the right of size one more than the set on the left. Then we ensure that distances between consecutive elements are doubled at each step and the first and last element are selected.

Finally, we will need the following MSO formulas $root^{(k)}(P_1, P_2)$ which checks if $|P_2| = |P_1|^k$, assuming k is a power of two; $div(P_1, P_2)$ which checks if $|P_1|$ divides $|P_2|$; and $mod_L(P_1, P_2, R)$ which is true if $|P_2| \bmod |P_1| = |R|$ and $|P_1| \leq L$.

4.2 Hardness for Unary Strings and Paths

Theorem 3. *Let f be an elementary function and c a constant. If there exists an algorithm which, given a unary string w of length n and an MSO formula ϕ decides if $w \models \phi$ in time $f(|\phi|)n^c$ then $EXP=NEXP$.*

Proof (sketch).

Suppose we are given a non-deterministic Turing machine running in time $T = 2^{n^k}$ for k a power of two, and therefore using at most T cells of its tape, when given n bits of input. We are also given the n bits of input, which we interpret as a binary number $I \leq 2^n$. We must decide if the machine accepts using the hypothetical model-checking algorithm for unary words.

We construct a path of size $T^2(2I+1)$. Using div we locate a sub-path of size I (finding the largest odd divisor of the path) and a sub-path of size T^2 , which we divide into sections of size T using $root$. These sections will represent snapshots of the tape during the machine's execution. We ask for a subset of the elements of the path that encodes the tape cells on which the machine writes 1. Now we need to check two things: first, that the bits at the beginning of the tape correspond to the input. This can be done with exp and some arithmetic on the size of I . Second, that the bits

selected encode a correct computation. This is done by checking all pairs of elements from consecutive snapshots that correspond to the same cell. These are identified using eq_L (their distance is exactly T). \square

Corollary 2. *Let f be an elementary function and c a constant. If there exists an algorithm which, given a path P on n vertices and an MSO formula ϕ decides if $P \models \phi$ in time $f(|\phi|)n^c$ then $EXP=NEXP$.*

Corollary 3. *Let f be an elementary function, c a constant, and \mathcal{C} a class of graphs closed under edge sub-divisions. If there exists an algorithm which, given a graph $G \in \mathcal{C}$ on n vertices and an MSO formula ϕ decides if $G \models \phi$ in time $f(|\phi|)n^c$ then $EXP=NEXP$. The same is true if \mathcal{C} is closed under induced subgraphs and, for all $d > 0$ contains a graph with diameter d .*

Finally, we can extend the ideas given above to obtain an alternative, self-contained proof of a result given in [3]: MSO_2 model-checking on cliques is not in XP, unless $EXP=NEXP$. In [3] this is proved under the equivalent assumption $P_1 \neq NP_1$ (the $P \neq NP$ assumption for unary languages). That proof relies on the work of Fagin on graph spectra [6].

Here we can simply reuse the ideas of Theorem 3 by observing two basic facts: first, with an appropriate MSO_2 formula we can select a set of edges in the given clique that induces a spanning path. Therefore, we can assume we have the same structure as in the case of paths. Second, the eq_L predicate can be constructed in constant size, since two disjoint sets of vertices are equal if and only if there exists a perfect matching between them in the clique (and this is MSO_2 -expressible).

Corollary 4. *If there exists an algorithm which, given a clique K_n on n vertices and an MSO_2 formula ϕ decides if $K_n \models \phi$ in $n^{f(|\phi|)}$, for any function f , then $EXP=NEXP$.*

5 Tree-Depth

In this section we give a lower bound result that applies to the model-checking algorithm for trees of bounded height given by Gajarský and Hliněný [9]. We recall here the main result:

Theorem 4 ([9]).

Let T be a rooted t -colored tree of height $h \geq 1$, and let ϕ be an MSO sentence with r quantifiers. Then $T \models \phi$ can be decided by an FPT algorithm in time $O\left(\exp^{(h+1)}\left(2^{h+5}r(t+r)\right) + |V(T)|\right)$.

Theorem 4 is the main algorithmic tool used to obtain the recent elementary model-checking algorithms for tree-depth and shrub-depth given in [9] and [11], since in both cases the strategy is to interpret the graph into a colored tree of bounded height.

The running time given in Theorem 4 is an elementary function of the formula ϕ , but non-elementary in the height of the tree. Though we would very much like to avoid that, it is not hard to see that the dependence on at least one of the parameters must be non-elementary, since allowing h to grow eventually gives the class of all trees so the lower bound result of Frick and Grohe should apply.

It is less obvious however what the height of the exponentiation tower has to be exactly, as a function of h , the height of the tree. The fact that we know that the height of the tower must be unbounded (so that we eventually get a non-elementary function) does not preclude an algorithm that runs in time $\exp^{(\sqrt{h})}(|\phi|)$ or, less ambitiously, $\exp^{(h/2)}(|\phi|)$, or even $\exp^{(h-5)}(|\phi|)$. Recall that we are trying to determine the number of levels of exponentiation in the running time here, so shaving off even an additive constant would be a non-negligible improvement.

We show that even such an improvement is probably impossible, and Theorem 4 determines precisely the complexity of MSO model-checking on colored trees of height h , at least in the sense that it gives exactly the correct level of exponentiations. We establish this fact assuming the ETH, by combining lower bound ideas which have appeared in [8] and [15]. More specifically, the main technical obstacle is comparing indices, or in other words, counting economically in our construction. For this, we use the tree representation of numbers of [8] pruned to height $h - 1$. We then use roughly $\log^{(h)} n$ colors to differentiate the leaves of the constructed trees.

Theorem 5. *If for some constant $h \geq 1$ there exists a model-checking algorithm for t -colored rooted trees of height h that runs in $\exp^{(h+1)}(o(t)) \cdot \text{poly}(n)$ time for trees with n vertices then the Exponential Time Hypothesis fails.*

6 Conclusions and Open Problems

We have proved non-elementary lower bounds for FO logic on cographs and MSO logic on uncolored paths. The hope is that, since these lower bounds concern very simple graph families, they can be used as “sanity checks” guiding the design of future graph widths. We have also given a lower bound for MSO logic on colored trees of bounded height. It would be interesting to see if this can be extended to uncolored trees.

Finally, let us mention that a promising direction in this area that we did not tackle here is that of alternative logics, besides FO and MSO variants. One example is the meta-theorems given by Pilipczuk [16] for a kind of modal logic. The algorithmic properties of such logics are still mostly unexplored but they may be a good way to evade the lower bounds given in [8] and this paper.

References

1. V. Chvátal and P. L. Hammer. Aggregation of inequalities in integer programming. *Annals of Discrete Mathematics*, 1:145–162, 1977.
2. B. Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.
3. B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.
4. B. Courcelle and S. Olariu. Upper bounds to the clique width of graphs. *Discrete Applied Mathematics*, 101(1-3):77–114, 2000.
5. A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Approximation schemes for first-order definable optimisation problems. In *LICS*, pages 411–420. IEEE Computer Society, 2006.
6. R. Fagin. Generalized first-order spectra and polynomial-time recognizable sets. *Proceedings American Mathematical Society*, 1974.
7. M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001.
8. M. Frick and M. Grohe. The complexity of first-order and monadic second-order logic revisited. *Ann. Pure Appl. Logic*, 130(1-3):3–31, 2004.
9. J. Gajarský and P. Hliněný. Faster deciding mso properties of trees of fixed height, and some consequences. In D. D’Souza, T. Kavitha, and J. Radhakrishnan, editors, *FSTTCS*, volume 18 of *LIPICs*, pages 112–123. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.
10. R. Ganian. Twin-cover: Beyond vertex cover in parameterized algorithmics. In D. Marx and P. Rossmanith, editors, *IPEC*, volume 7112 of *Lecture Notes in Computer Science*, pages 259–271. Springer, 2011.
11. R. Ganian, P. Hliněný, J. Nešetřil, J. Obdržálek, P. O. de Mendez, and R. Ramadurai. When trees grow low: Shrubs and fast mso1. In B. Rovan, V. Sassone, and P. Widmayer, editors, *MFCS*, volume 7464 of *Lecture Notes in Computer Science*, pages 419–430. Springer, 2012.
12. M. Grohe. Logic, graphs, and algorithms. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(091), 2007.
13. P. Hliněný, S. il Oum, D. Seese, and G. Gottlob. Width parameters beyond tree-width and their applications. *Comput. J.*, 51(3):326–362, 2008.
14. S. Kreutzer. On the parameterized intractability of monadic second-order logic. *Logical Methods in Computer Science*, 8(1), 2012.
15. M. Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012.
16. M. Pilipczuk. Problems parameterized by treewidth tractable in single exponential time: A logical approach. In F. Murlak and P. Sankowski, editors, *MFCS*, volume 6907 of *Lecture Notes in Computer Science*, pages 520–531. Springer, 2011.

A Omitted Material

A.1 Proof of Theorem 2

Proof. Suppose that we are given a binary word w and an FO formula ϕ . We will reduce the problem of deciding if $w \models \phi$ to the problem of deciding whether $G \models \phi'$ for a threshold graph G and a FO formula ϕ' which we will construct.

First, let us describe G , and since it's a threshold graph we can describe it as a string over the alphabet $\{u, j\}$. The graph G starts with uu . Then, for each character of w , if it is a 0 we append uj to the description of G , otherwise we append ujj . So, for example the graph corresponding to $w = 01101$ would have description $uujuujjujjjujj$. Notice that, since the last character in the description is a j , the graph is connected.

Now we need to interpret the formula ϕ into the new context. To do this, let's first observe some basic properties of our graph. First, a vertex in this graph is a union vertex if and only if its neighborhood is a clique. To see this, note that union vertices are only connected to join vertices, which form a clique. All join vertices on the other hand are connected to the first two union vertices which are not connected. Second, all union vertices, except the first two (dummy) vertices have at least one join vertex as a non-neighbor, namely at least the first join vertex.

We thus define the following formulas

$$\begin{aligned} \text{union}(x) &:= \forall y \forall z (E(x, y) \wedge E(x, z) \wedge y \neq z) \rightarrow E(y, z) \\ \text{main}(x) &:= \text{union}(x) \wedge (\exists y \neg \text{union}(y) \wedge \neg E(x, y)) \end{aligned}$$

This will allow us to simulate selecting a character in the word by selecting the union vertex which represents the corresponding pair or triple of vertices in the graph.

Now we also need to encode the \prec and P_1 predicates. We define

$$\begin{aligned} \text{prec}(x, y) &:= \exists z \neg \text{union}(z) \wedge E(x, z) \wedge \neg E(y, z) \\ \text{one}(x) &:= \exists y \exists z (y \neq z) \wedge \neg \text{union}(y) \wedge \neg \text{union}(z) \wedge E(x, y) \wedge E(x, z) \\ &\quad \wedge \forall w (\text{main}(w) \wedge \text{prec}(x, w)) \rightarrow (\neg E(y, w) \wedge \neg E(z, w)) \end{aligned}$$

The intuition for the first is that, if x, y are two union vertices that represent two different blocks, x precedes y if and only if there exists some join vertex connected to x but not y . For P_1 we have that x is a union

vertex representing a *ujj* block if and only if there exist two join vertices connected to it and not connected to any union vertex that comes later in the description.

Given the above it is straightforward to produce the formula ϕ' from ϕ : the formulas *prec* and *one* are used to translate the corresponding atomic predicates \prec and P_1 , while we inductively replace $\exists x\psi(x)$ with $\exists x \text{main}(x) \wedge \psi'(x)$ where $\psi'(x)$ is the translation of $\psi(x)$. It is not hard to see that $|\phi'| = O(\phi)$ while the order of G is $O(|w|)$. \square

A.2 Proof of Corollary 1

Proof. It suffices to prove this if the class is closed under union and join, because if it's closed under union and complement we get closure under join “for free”. The proof is immediate if $K_1 \in \mathcal{C}$, since then the class contains threshold graphs. Otherwise, let G_m be the graph of the smallest order in the class and say it has k vertices. We will construct a graph as in the proof of Theorem 2, except that for each vertex we would be adding in that case we will add a copy of G_m . More specifically, for each union vertex of the threshold graph we add a disconnected copy of G_m to the graph we are constructing and for each join vertex we add a copy of G_m and connect all its vertices to all previously added vertices. It's easy to see that the graph we have constructed is still in \mathcal{C} .

It is now not hard to see how to translate the proof of Theorem 2 in this case. Replace every $\exists x\phi(x)$ with $\exists x_1\exists x_2\dots\exists x_k G_m(x_1, \dots, x_k) \wedge \phi(x_1)$, where G_m is a formula stating that the x_i 's have the structure of G_m (that is, they are all distinct and have the same edges as G_m) and they all have the same neighbors in the rest of the graph. Knowing that the x_i 's form a copy of G_m that corresponds to a single character, we then take one representative and use it in the rest of the formula (note that if we take one representative from each copy the result is a threshold graph).

This trick is sufficient to translate the formulas for *main* and *prec*. The only place where we may run into a problem are the *union* and *one* formula, because we use the \neq predicate there. Since we are picking x_1 as an arbitrary representative of a copy of G_m , if G_m has a non-trivial automorphism it could be the case that the k vertices that correspond to $\exists y$ and the k vertices that correspond to $\exists z$ are assigned to the same copy of G_m , but $y_1 \neq z_1$. To avoid this case we just need to add an extra formula after the quantification of y, z stating that all y_i, z_j are pairwise distinct. \square

A.3 Formal definition of eq_L

Let us now give a formal definition of eq_L . First, we need to be able to recognize sections. Assume that we have a set of elements U and a subset $P \subseteq U$. As explained P divides U into sections so we define a formula $section(S, U, P)$ that will be true if S is such a section.

$$\begin{aligned}
consec(S, U) &:= \forall x \forall y \forall z \left((x \in S) \wedge (y \in S) \wedge (z \in U) \wedge (x \prec z) \right. \\
&\quad \left. \wedge (z \prec y) \right) \rightarrow z \in S \\
partsection(S, U, P) &:= S \subseteq U \wedge consec(S) \wedge \forall x, y \in S (x \in P \leftrightarrow y \in P) \\
section(S, U, P) &:= partsection(S, U, P) \wedge \\
&\quad \forall S' (S \subseteq S' \wedge partsection(S', U, P) \rightarrow S' = S)
\end{aligned}$$

Informally, $consec$ checks if S is a contiguous subset of U . Then, S is a partial section if it's a subset of either P or $U \setminus P$ and it represents a contiguous subset of U elements. S is a section if it's a maximal partial section.

Assuming two sets S_1, S_2 represent consecutive sections and we have a set B which is supposed to encode a binary number i in S_1 and $i + 1$ in S_2 we check this with the following formula:

$$\begin{aligned}
next(S_1, S_2, B) &:= \exists S_1^L \exists s^1 \exists S_1^R \exists S_2^L \exists s^2 \exists S_2^R \\
&\quad \bigwedge_{i=1,2} \left(S_i^L \subseteq S_i \wedge s^i \in S_i \wedge S_i^R \subseteq S_i \right) \wedge \\
&\quad \bigwedge_{i=1,2} \left(\forall x \in S_i^L \forall y \in S_i^R (x \prec s^i \wedge s^i \prec y) \right) \wedge \\
&\quad \bigwedge_{i=1,2} \left(S_i \subseteq S_i^L \cup \{s^i\} \cup S_i^R \right) \\
&\quad \wedge s^2 \in B \wedge s^1 \notin B \wedge (S_1^R \subseteq B) \wedge (S_2^R \cap B = \emptyset) \\
&\quad \wedge eq_L(S_1^R, S_2^R) \\
&\quad \wedge same(S_1^L, S_2^L, B, B) \\
same(S_1, S_2, B_1, B_2) &:= \exists f_1 \exists f_2 \bigwedge_{i=1,2} (\forall x \in S_i (f_i \prec x \vee f_i = x)) \wedge \\
&\quad \forall S_1^L \forall S_2^L \left(\bigwedge_{i=1,2} (f_i \in S_i^L) \right. \\
&\quad \left. \wedge consec(S_i^L, S_i) \wedge eq_L(S_1^L, S_2^L) \right) \rightarrow
\end{aligned}$$

$$\begin{aligned} & \exists l_1 \exists l_2 \bigwedge_{i=1,2} \left((\forall x \in S_i^L \ x \prec l_i \vee l_i = x) \right. \\ & \left. \wedge (l_1 \in B_1 \leftrightarrow l_2 \in B_2) \right) \end{aligned}$$

Informally, *next* partitions the sets S_1, S_2 into a left and right part and identifies a vertex in the middle. Assuming $|S_1| = |S_2|$ respective parts have the same size; the right part of S_1 corresponds to all 1 digits and the right part of S_2 to all 0. The left parts have to encode the same number, which is checked by *same*. The idea here is that if we select equal length contiguous prefixes of the sets we are checking the last element will encode the same digit in both. We allow *same* to use different sets B_1, B_2 to read the encoding in S_1, S_2 . This extra generality will be useful when we reuse this formula later.

We will use *next* only for neighboring sections. To check if two disjoint sections are indeed adjacent we define:

$$adj(S_1, S_2, U) := \exists x \in S_1 \exists y \in S_2 \forall z \in U (z \prec x \vee y \prec z \vee z = x \vee z = y)$$

Informally, *adj* is true if S_1 is a section that directly precedes the section S_2 , because then x is the last element of S_1 and y is the first element of S_2 and there are no elements between them.

We are now ready to define eq'_L for $L' = L2^L$.

$$\begin{aligned} eq'_L(P_1, P_2) & := \exists Q_1 \exists Q_2 \exists R_1 \exists R_2 \exists B_1 \exists B_2 \\ & \bigwedge_{i=1,2} (Q_i \subseteq P_i \wedge R_i \subseteq P_i \wedge B_i \subseteq P_i) \wedge eq_L(R_1, R_2) \wedge \\ & \left(\forall S_1 \forall S_2 \left(section(S_1, P_1 \setminus R_1, Q_1) \wedge \right. \right. \\ & \left. \left. section(S_2, P_2 \setminus R_2, Q_2) \right) \rightarrow eq_L(S_1, S_2) \right) \wedge \\ & \bigwedge_{i=1,2} \left(\forall S \forall S' ((S \cap S' = \emptyset) \wedge section(S, P_i \setminus R_i, Q_i) \wedge \right. \\ & \left. section(S', P_i \setminus R_i, Q_i) \wedge adj(S, S', P_i \setminus R_i) \right. \\ & \left. \rightarrow next(S, S', B_i) \right) \wedge \\ & \exists S_1^F \exists S_2^F same(S_1^F, S_2^F, B_1, B_2) \wedge \\ & \bigwedge_{i=1,2} \left(section(S_i^F, P_i \setminus R_i, Q_i) \wedge \right. \\ & \left. \forall S' (section(S', P_i \setminus R_i, Q_i) \wedge (S' \cap S_i^F = \emptyset)) \right) \end{aligned}$$

$$\begin{aligned}
& \rightarrow \exists x \in S' \exists y \in S_i^F x \prec y) \wedge \\
& \forall S \bigwedge_{i=1,2} \text{section}(S, P_i \setminus R_i, Q_i) \rightarrow (S \setminus B_i \neq \emptyset) \wedge \\
& \exists S \bigwedge_{i=1,2} \text{section}(S, P_i \setminus R_i, Q_i) \wedge (S \cap B_i = \emptyset)
\end{aligned}$$

This is rather long, so let us explain it intuitively. We want to test if $|P_1| = |P_2|$, so we demand the following:

- From both we remove a remainder set R_i , and we make sure that the remainder sets are equal.
- We use Q_1, Q_2 to partition the two sets into sections. All sections of the first set must be equal in size to all sections of the second (therefore, all sections in both sets are equal).
- Select the sets B_i which will encode binary numbers in the sections. For each two disjoint sections which are consecutive check that they encode consecutive numbers.
- Find the last section on each set (S_1^F, S_2^F). Check that they encode the same number.
- Check that no section encodes a number made up only of 1s, so we don't have an overflow in the counting.
- Check that the first section on each set encodes the number zero (that is, it has no elements from B_i). For this it's sufficient to check that some section encodes zero, since we have already established proper ordering.

Proof (Lemma 1).

Correctness follows by induction and the definition of the construction given above. For the size bound, note that for $L' = L2^L$ we have $|eq_{L'}| = O(|eq_L|)$, since the definition of $eq_{L'}$ uses eq_L a constant number of times. It follows that there exists a constant c such that for all k we have $|eq_{exp(k)(1)}| = O(c^k)$. The result follows by setting $k = \log^* L$. \square

A.4 Formal Definition of *exp* and Other Arithmetic

Again, we first define some auxilliary formulas. Checking if a set is twice as large as another can be done as follows:

$$\text{double}(S_1, S_2) := \exists S'(S' \subseteq S_2) \wedge eq_L(S_1, S') \wedge eq_L(S_1, S_2 \setminus S')$$

If we are given three elements x, y, z such that $x \prec y \prec z$ we can check that the distance is doubled as follows:

$$\begin{aligned} ddist(x, y, z, U) := & \exists S_1 \subseteq U \wedge \forall u ((x \prec u \vee x = u) \wedge (u \prec y)) \leftrightarrow u \in S_1 \\ & \wedge \exists S_2 \subseteq U \wedge \forall u ((y \prec u \vee y = u) \wedge (u \prec z)) \leftrightarrow u \in S_2 \\ & \wedge double(S_1, S_2) \end{aligned}$$

Informally, we select the sets S_1, S_2 as the sets of elements starting from x and up to (but not including) y , and starting from y and up to (but not including) z . The second set must be twice as large. Now we can define *exp*:

$$\begin{aligned} exp(P_1, P_2) := & \exists Q \exists f \exists s \exists l (Q \subseteq P_2) \wedge f \in Q \wedge s \in Q \wedge l \in Q \wedge f \prec s \wedge \\ & \forall u (u \in P_2) \rightarrow (u \prec l \vee u = l) \wedge (s \prec u \vee s = u \vee f = u) \\ & eq_L(P_1, Q \setminus \{l\}) \wedge \\ & \forall x \forall y \forall z (x \in Q) \wedge (y \in Q) \wedge (z \in Q) \wedge (x \prec y) \wedge (y \prec z) \wedge \\ & \neg \exists u (u \in Q) \wedge ((x \prec u) \wedge (u \prec y)) \vee ((y \prec u) \wedge (u \prec z)) \\ & \rightarrow ddist(x, y, z, P_2) \end{aligned}$$

In words, we select a set Q from P_2 so that any three consecutive elements double consecutive distances. The total size of P_2 must then be $(1 + 2 + 4 + \dots + 2^{|P_1|-1}) + 1 = 2^{|P_1|}$, where the sum is obtained by adding the consecutive distances, and we add one at the end because the last element l was not counted.

Finally, we would like a formula that calculates the k -th root, assuming (to keep things simple) that k is a power of two. We will first need a formula $root_L(P_1, P_2)$ which will simply be true if $|P_2| = |P_1|^2$.

$$\begin{aligned} root_L(P_1, P_2) := & \exists Q (Q \subseteq P_2) \wedge (\forall S section(S, P_2, Q) \rightarrow eq_L(S, P_1)) \wedge \\ & \exists S' eq_L(S', P_1) \wedge \forall S section(S, P_2, Q) \rightarrow (|S \cap S'| = 1) \end{aligned}$$

Informally, we can divide P_2 into sections of size $|P_1|$ and if we select a set S' that contains exactly one representative from each section then $|S'| = |P_1|$. Given this formula, we can now define $root_L^{(k)}(P_1, P_2)$, the formula that calculates k -th roots as follows:

$$\begin{aligned} root_L^{(2)}(P_1, P_2) & := root_L(P_1, P_2) \\ root_L^{(k)}(P_1, P_2) & := \exists Q root_L(Q, P_2) \wedge root_L^{(k/2)}(P_1, Q) \end{aligned}$$

It is not hard to see that $root_L^{(k)}(P_1, P_2)$ is true if $|P_2| = |P_1|^k$, assuming that k is a power of two.

We will need to perform division:

$$\begin{aligned} div(P_1, P_2) &:= \exists Q \subseteq P_2 \forall S section(S, P_2, Q) \rightarrow eq_L(S, P_1) \\ less_L(P_1, P_2) &:= \exists S (P_1 \subset S) \wedge eq_L(S, P_2) \\ mod_L(P_1, P_2, R) &:= \exists S \subseteq P_2 \wedge eq_L(S, R) \wedge div(P_1, P_2 \setminus S) \wedge less(R, P_1) \end{aligned}$$

The div formula decides if $|P_1|$ exactly divides $|P_2|$ by partitioning P_2 into sections of size $|P_1|$. Using this we can then calculate remainders.

A.5 Proof of Theorem 3

Proof. Suppose that we are given a non-deterministic Turing machine that runs in time 2^{n^k} , for some constant k , when given n bits of input. We will use the hypothetical algorithm to predict whether the machine accepts an arbitrary input in deterministic exponential time.

Let us discuss some technical details about the machine. Without loss of generality, assume that k is a power of two and we are given a non-deterministic machine that always terminates in time at most $T = 2^{n^k}$. Assume that the machine uses a binary alphabet, and without loss of generality it never uses more than T cells of tape. Also without loss of generality, we may assume that the first thing the machine does is non-deterministically guess a string of bits and use it to fill out its tape. From that point on the machine behaves deterministically, that is, there is a finite set of states Q and a transition function $\delta : Q \times \{0, 1\} \rightarrow Q \times \{0, 1\} \times \{L, S, R\}$, that tells the machine for each state and cell character, which state to go to next, what to write on the current cell and whether to move left, right, or stay at the same cell. The state set Q contains a special state q_{acc} such that if the machine ever enters this state it automatically accepts and never leaves this state.

Suppose that we have been given the description of such a machine with $|Q|$ states, where $|Q|$ is independent of the input, and n bits of input. We will construct a unary string w of appropriate length and an MSO formula ϕ such that $w \models \phi$ if and only if the machine would accept this input.

Let I be the number whose binary representation is exactly the input given to the machine (so $I \leq 2^n$). Construct a unary string w of length $L = (2I+1)T^2$, where we recall that T is the upper bound on the machine's running time. Now we need to construct the formula ϕ .

Rather than giving all formal details, we will now give a high-level description of ϕ and the reader may verify that ϕ can indeed be constructed with the tools from the previous section. Our formula will first ensure the following:

- First, it will identify a subset of the input with size I and another with size T^2 . This is achievable by observing that the largest odd divisor of L is $2I + 1$, so we simply ask for the largest odd set whose size exactly divides the input.
- Using the root formula we partition the set of size T^2 into T equal-size sections. Each will correspond to a snapshot of the tape during a step in the machine’s execution.
- Identify the first section of the tape, and then identify a prefix of it with size n^k (this can be done with the *exp* formula). Identify a prefix of that with size n (this can be done with the *root*^(k) formula, since k is a constant and a power of two). This is where the machine’s input will initially stored.

It should be clear that the above can be expressed in MSO with the formulas of the previous section. So at this point, we have identified T sections, each of size T , to represent the machine’s tape. Each element thus naturally corresponds to a specific cell at a specific point in time during the machine’s run. We have also identified a special part at the start of the first tape section with size n , where we will check that the input is stored, and a set of length I that encodes the input. Now we ask for the existence of a subset B of elements that will indicate the cells of the tape where 1 is written. We also ask for the existence of $|Q|$ sets, call them H_i , $i \in Q$. The intended meaning is that if a certain element from one of the tape sections is in H_i , then the machine was in state i at the point in time corresponding to that section and the machine’s head was located at the cell corresponding to that element.

Once the above sets have been selected all of the machine’s computation has been encoded. Now we just need to check that it’s correct and accepting. We thus express the following conditions:

- Ensure the input is correctly encoded at the start of the tape. To check the bit at position i we observe that the contiguous subset of the tape from the beginning to that bit has size i . Using *exp* we can construct a set of size 2^i and then a set of size 2^{i+1} . We “calculate” the remainder of the set of length I with the set of size 2^{i+1} . If the result has size at least 2^i the bit has to be 1 otherwise 0. With this idea we can check the correctness of all input bits.

- Ensure that the machine transitions correctly. We look at pairs of elements that correspond to the same tape cell in consecutive steps in time, that is, tape elements whose distance is exactly T , which can be verified with the eq_L formula. If the first has no H_i label then either both have B or neither does. If the first has an H_i label we check that the B label changes appropriately for the other and an H_j label is used appropriately for the other or one of its neighbors, depending on the transition function.
- Finally, check that in each section of the tape exactly one element has an H_i label, and that it has exactly one. Also, check that some element eventually gets the H_{qacc} label.

All the above requirements can be checked with an MSO formula with constant size (assuming $k, |Q|$ constant), except for uses of the eq_L predicate, which has size $2^{O(\log^* L)}$. So the whole formula also has size $2^{O(\log^* L)}$ and $L = 2^{poly(n)}$ so $|\phi| = 2^{O(\log^* n)}$.

Suppose that an algorithm with running time $f(|\phi|)|w|^c$ existed for elementary f . Then, there exists d such that $f(x) \leq \exp^{(d)}(x)$. So the running time is at most $\exp^{(d+1)}(O(\log^* n))2^{poly(n)} = 2^{poly(n)}$. \square

A.6 Proof of Corollary 3

Proof. If the class is closed under induced subgraphs and for each $d > 0$ there is a graph in the class with diameter d , then the class contains all paths and therefore we can invoke Corollary 2. To see this, for each d take the graph with diameter d and let u, v be two vertices with shortest path distance d . The graph induced by u, v and the vertices that make up a shortest path from u to v is a path, since if more edges were induced a shorter path would exist from u to v . Thus, the class contains a path with $d + 1$ vertices.

If the class is closed under edge subdivisions we can reduce (in fact interpret) the MSO model checking problem on paths to MSO model checking on the class. We have a path with n vertices and by the proof of Theorem 3 we can assume n to be even. Select the smallest graph in \mathcal{C} , call it G_m , and subdivide its edges an appropriate number of times so that all maximal connected sets of degree two vertices have odd size. Select one such set and subdivide its edges so that it has size n (this is always possible if we started with a sufficiently large n). It is not hard to amend the original formula so that it first locates this path of size n that we created (it's now the only maximal connected set of degree two

vertices with even size) and only works with vertices from it. Since we only subdivided edges the graph we have is still in \mathcal{C} . \square

A.7 Proof of Corollary 4

Proof. The proof follows similar lines as in Theorem 3, so we explain here the differences. First, we must implement the \prec predicate on the clique. This is achieved by selecting a set of edges that induces a path and then using the same tricks as we use to simulate MSO for strings with MSO_1 for paths. Second, we must implement an $eq(P_1, P_2)$ predicate with constant size. If we do both of these, the rest of the proof of Theorem 3 goes through unchanged, since the formula we construct only uses the \prec and $=$ predicates and has constant size except for the eq_L predicate.

The main observation is that we can ask for the existence of a perfect matching between two disjoint sets as follows (we denote by \exists_e edge-set quantification, and $I(x, e)$ the vertex-edge incidence relation):

$$\begin{aligned} pm(S_1, S_2) &:= \exists_e F \forall x (x \in S_1 \vee x \in S_2) \rightarrow \\ &\quad \exists e \in F I(x, e) \wedge (\forall f \in F I(x, f) \rightarrow e = f) \\ &\quad \wedge \forall e \in F \exists x \exists y (x \in S_1) \wedge (y \in S_2) \wedge I(x, e) \wedge I(y, e) \end{aligned}$$

Informally, we ask if there exists a set of edges so that each vertex has exactly one incident edge and all edges have one endpoint in each set. It should be clear that there exists a perfect matching between two disjoint sets if and only if the sets have the same size. Thus, we can define eq as follows:

$$eq(P_1, P_2) := pm(P_1 \setminus P_2, P_2 \setminus P_1)$$

Thus, all that is left is to implement the \prec predicate. We give here a high-level argument. First, we will ask for the existence of a set F of edges with the following properties:

- All vertices have exactly two edges of F incident on them.
- For any partition of the vertices there exists an edge from F with endpoints on both sides (connectivity).

It is not hard to see that F induces a spanning cycle. Let F' be the set obtained from F by removing an arbitrary edge and let s be one of

the endpoints of the removed edge. We will know say that $x \prec y$ if any subset of edges from F' that connects s to y must touch x .

The size of the clique we construct is the same as the length of the string in Theorem 3. So if there exists an algorithm running in time polynomial in the order of the input clique for fixed-size formulas then we can simulate a NEXP Turing machine in exponential time and $\text{EXP}=\text{NEXP}$. \square

A.8 Proof of Theorem 5




Fig. 3. An example of the graphs constructed in the proof of Theorem 5. Assume we have two labels available, one represented by solid grey fillings and the other with a dashed box around vertices that have it. The first 4 numbers ($0, \dots, 3$) can be represented by a single vertex. Numbers up to $2^4 - 1$ can be represented with trees of height 1, numbers up to $2^{2^4} - 1$ with trees of height 2, etc.

Proof. As usual in such proofs, the main obstacle is how to encode numbers up to n economically in terms of the height of the constructed tree and the colors used. Fix some $h \geq 2$ (we will handle the case $h = 1$ in the end). We have at our disposal around $\log^{(h)} n$ colors. By using them we can create $2^{\log^{(h)} n} = \log^{(h-1)} n$ vertices which we can distinguish by using a different set of colors on each vertex. To go from there to n we will use the trick of [8] which, roughly speaking, gives exponentially more counting power with each level of height added. Thus, we will manage to represent numbers up to n with trees of height $h - 1$.

Let us now be more precise. We have at our disposal $\log^{(h)} n$ colors, number them $0, \dots, \log^{(h)} n - 1$. We will define for each $i \in \{0, \dots, n - 1\}$ a rooted colored tree T_i . The construction is inductive:

- If $i \in \{0, \dots, \log^{(h-1)} n - 1\}$ then i has a binary representation with at most $\log^{(h)} n$ bits, say $b_k b_{k-1} \dots b_1 b_0$ with $k \leq \log^{(h)} n - 1$. The tree T_i is a single vertex colored with exactly the colors j such that $b_j = 1$.

- Suppose that we have defined T_i for $i \in \{0, \dots, \log^{(k)} n - 1\}$ for some $k \geq 1$. We will now define T_i for $\log^{(k)} n \leq i \leq \log^{(k-1)} n - 1$. As previously, write down the binary representation of i , which has at most $\log^{(k)} n$ bits. For each j such that $b_j = 1$ construct a copy of T_j (we already know how to do this by the inductive hypothesis). Add a new vertex, which will be the root of the new tree, and connect it to the roots of the constructed trees.

Some examples of the above construction are given in Figure 3. Now let i be an integer such that $1 \leq i \leq h$. We observe that the above construction represents numbers which are at most $\leq \log^{(h-i)} n - 1$ with trees of height $i - 1$. This can be proved by induction: for $i = 1$ trees of height 0 (that is, single vertices) are used to represent numbers up to $\log^{(h-1)} n - 1$. For the inductive case, notice that each level of height added increases the maximum number representable exponentially. As a result, the numbers $0, \dots, n - 1$ can be represented with a tree of height $h - 1$.

We can now also define an $eq_k(x, y)$ predicate, that will be true if and only if x, y are the roots of two trees of height at most k representing the same number. Again we proceed inductively:

- It's easy to define a simple propositional predicate $\text{samecols}(x, y)$. The predicate will be true if x, y have exactly the same colors from the set $\{0, 1, \dots, \log^{(h)} n - 1\}$. Using this, for $k = 0$ we set $eq_0(x, y) := \text{samecols}(x, y) \wedge \forall z \neg C(z, x) \wedge \neg C(z, y)$. In other words, x, y are equal if they have the same colors and no children.
- Suppose $eq_k(x, y)$ is defined, we will define $eq_{k+1}(x, y)$. We set

$$eq_{k+1}(x, y) := \text{samecols}(x, y) \wedge \forall u (C(u, x) \vee C(u, y)) \rightarrow \\ \exists v \left(eq_k(u, v) \wedge \right. \\ \left. (C(u, x) \rightarrow C(v, y)) \wedge (C(u, y) \rightarrow C(v, x)) \right)$$

In words, x, y have the same colors and for every vertex that is the child of one of them there exists a vertex that is a child of the other and these two vertices represent the same number.

It is not hard to see that the formula $eq_h(x, y)$ uses $O(h)$ quantifiers. We are now ready to describe our construction.

Fix $h > 1$ and start with an instance of 3SAT with n variables and suppose that these variables are named $x_i, i \in \{n, \dots, 2n - 1\}$. The reason we number the variables this way is that it will be convenient for all of

them to have an index high enough that a non-trivial tree is needed to describe it. For each variable x_i construct a copy of the tree T_i as described above. We will make use of $\log^{(h)} 2n = \log^{(h)} n + o(1)$ colors so trees have height at most $h - 1$. Color the roots of all these trees with a new color, call it v .

For each clause $(l_i \vee l_j \vee l_k)$ where l_i, l_j, l_k are literals, that is, positive or negative appearances of the variables x_i, x_j, x_k respectively, construct three trees T_i, T_j, T_k . Introduce six new colors, call them $c_{p,q}$ for $p \in [3], q \in \{0, 1\}$. If the literal l_i is positive then color the children of the root of the tree T_i with $c_{1,1}$, otherwise color them with $c_{1,0}$. Similarly, color the children of the root of T_j with $c_{2,1}$ if l_j is positive and $c_{2,0}$ if it's negative and the children of the root of T_k with $c_{3,1}$ or $c_{3,0}$. Notice that we know that all trees have height at least 1 because $h \geq 2$ and variables are numbered $n, \dots, 2n - 1$, so all roots do have children. Finally, merge the roots of T_i, T_j, T_k into a single vertex. We introduce a new color, call it c , and use it to color the new root of the tree that represents each clause. To complete the construction, add a new root vertex to the graph and make all roots of previously constructed trees its children. This creates a tree with height h . The root has one v -colored child for each variable and one c -colored child for each clause of the 3SAT formula.

Now, let us describe a formula with $O(h)$ quantifiers that will check if the original 3SAT instance was satisfiable. Informally, we will ask if there exists a set of variables, represented by a subset of the vertices colored with v , such that setting these to true and the rest to false satisfies the formula. To do this, we need to be able to check if a variable appears positive or negative in a clause. Let's define a predicate $pos_i(x, y)$ which will be true if variable x appears positive in position i (where $i \in [3]$) in the clause represented by y (so x is assumed to be the root of a variable tree and y is assumed to be the root of a clause tree).

$$\begin{aligned} pos_i(x, y) := & \forall u \left(C(u, x) \vee (C(u, y) \wedge P_{c_{i,1}}(u)) \right) \rightarrow \\ & \exists v \left(eq_h(u, v) \wedge \right. \\ & \left. (C(u, x) \rightarrow (C(v, y) \wedge P_{c_{i,1}}(y))) \wedge (C(u, y) \rightarrow C(v, x)) \right) \end{aligned}$$

The logic here is exactly the same as in the eq predicate, except that we only take into account the children of the clause node that correspond to the i -th literal. It's easy to see how to make a similar predicate neg_i for negative appearances (change $P_{c_{i,1}}$ to $P_{c_{i,0}}$). Given these, the complete formula is:

$$\begin{aligned}
isSAT := & \exists S(\forall x(x \in S \rightarrow P_v(x))) \wedge \\
& \forall y(P_c(y) \rightarrow \exists z(P_v(Z) \wedge \\
& (\bigvee_{i \in [3]}(pos_i(z, y) \wedge z \in S)) \vee (\bigvee_{i \in [3]}(neg_i(z, y) \wedge \neg z \in S))))
\end{aligned}$$

In words, there exists a set of variables S (which will be set to true), such that for each clause there exists a variable appearing in it that satisfies it, that is, it belongs in S if and only if its appearance is positive.

The formula has h quantifiers, we have used $\log^{(h)} n + O(1)$ colors and have constructed a tree of height h and size polynomial in n . If there exists a model-checking algorithm for t -colored trees running in $\exp^{(h+1)}(o(t))|V|^c$ this gives a $2^{o(n)}$ algorithm for 3-SAT.

The only thing left is the case $h = 1$. Here each variable and each clause will be represented by a single vertex and, since we have $O(\log n)$ colors available, the colors alone will be sufficient to compare indices. It's not hard to see how to encode the whole structure of the formula using $7 \log n$ colors. The first $\log n$ colors are used for the variable vertices. Then we need 6 sets of $\log n$ distinct colors to encode the appearances of literals into the clauses, for each combination of position and positivity. This makes it straightforward to implement pos_i and neg_i by comparing appropriate sets of colors on the two vertices. The rest of the formula is unchanged. \square