Lower Bounds on Dynamic Programming for Maximum Weight Independent Set

Tuukka Korhonen ✉ 🏠 📵
Department of Computer Science, University of Helsinki, Finland

—— Abstract ——

We prove lower bounds on pure dynamic programming algorithms for maximum weight independent set (MWIS). We model such algorithms as tropical circuits, i.e., circuits that compute with max and + operations. For a graph G, an MWIS-circuit of G is a tropical circuit whose inputs correspond to vertices of G and which computes the weight of a maximum weight independent set of G for any assignment of weights to the inputs. We show that if G has treewidth w and maximum degree d, then any MWIS-circuit of G has $2^{\Omega(w/d)}$ gates and that if G is planar, or more generally H-minor-free for any fixed graph H, then any MWIS-circuit of G has $2^{\Omega(w)}$ gates. An MWIS-formula is an MWIS-circuit where each gate has fan-out at most one. We show that if G has treedepth t and maximum degree d, then any MWIS-formula of G has $2^{\Omega(t/d)}$ gates. It follows that treewidth characterizes optimal MWIS-circuits up to polynomials for *all* bounded degree graphs and H-minor-free graphs, and treedepth characterizes optimal MWIS-formulas up to polynomials for all bounded degree graphs.

1 Introduction

In this paper, we prove lower bounds for *tropical circuits* computing the weight of a maximum weight independent set (MWIS) of a graph. A tropical circuit is a circuit with Max and Plus operations as gates. In particular, we consider *MWIS-circuits* of graphs. An MWIS-circuit of a graph G is a tropical circuit whose inputs correspond to the vertices of G and which computes the weight of a maximum weight independent set of G for any assignment of weights to the inputs. An MWIS-formula is an MWIS-circuit where each gate has fan-out at most one.

Our motivation for proving lower bounds for MWIS-circuits is that many algorithmic techniques for maximum weight independent set implicitly build an MWIS-circuit of the input graph, and therefore the running time of any algorithm resulting from such a technique is bounded from below by the minimum size of an MWIS-circuit of the graph. Examples of algorithmic techniques that build MWIS-circuits are dynamic programming over different kinds of decompositions of graphs [3, 8, 16] and dynamic programming over potential maximal cliques [10, 17, 29]. Examples of algorithmic techniques that build MWIS-formulas are branching [20, 34] and maximal independent set enumeration [32].

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).
Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 87; pp. 87:1–87:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1.1 Our Results

We prove unconditional lower bounds for sizes of MWIS-circuits and MWIS-formulas parameterized by graph parameters treewidth and treedepth, respectively. The lower bounds are exponential in treewidth and treedepth, and therefore well-known algorithms yield matching upper bounds for them [3, 18]. We emphasize that our lower bounds are not worst-case bounds over graph classes, but instead hold for each individual graph.

MWIS-Circuits and Treewidth

First we characterize optimal MWIS-circuits of bounded degree graphs.

▶ **Theorem 1.** *Let G be any graph with treewidth w and maximum degree d. Any MWIS-circuit of G has $2^{\Omega(w/d)}$ gates.*

Theorem 1 is optimal up to a factor d in the sense that for each pair w, d we can construct a graph with treewidth $\Omega(w)$ and maximum degree $O(d)$ that admits an MWIS-formula with $d2^{w/d}$ gates.

Then we extend the result to some graphs that may have high-degree vertices. A graph H is a *minor* of a graph G if it can be obtained from G by vertex deletions, edge deletions, and edge contractions. If H can be obtained by only vertex deletions and edge contractions, then it is an *induced minor*.

▶ **Theorem 2.** *Let G be any graph that contains an induced minor with treewidth w and maximum degree d. Any MWIS-circuit of G has $2^{\Omega(w/(d4^d))}$ gates.*

In Theorem 2 it is essential to require an induced minor instead of a minor because a complete graph with n vertices admits an MWIS-circuit of size $O(n)$, but contains all n-vertex graphs as minors.

A graph G is H-minor-free if it does not contain the graph H as a minor.

▶ **Corollary 3.** *Let H be any fixed graph and G any H-minor-free graph with treewidth w. Any MWIS-circuit of G has $2^{\Omega(w)}$ gates.*

Proof. For any fixed H, every H-minor-free graph of treewidth w contains an $\Omega(w) \times \Omega(w)$-grid as an induced minor [13][1]. An $\Omega(w) \times \Omega(w)$-grid has treewidth $\Omega(w)$ and maximum degree 4, so the result follows from Theorem 2. ◀

Corollary 3 implies a $2^{\Omega(w)}$ lower bound for all planar graphs because planar graphs are K_5-minor-free [27].

The following corollary follows from Theorem 1, Corollary 3, constant-factor treewidth approximation in $2^{O(w)}n^{O(1)}$ time [36], and dynamic programming over a tree decomposition [3].

▶ **Corollary 4.** *There is an algorithm which, given a bounded degree or H-minor-free graph G whose smallest MWIS-circuit has τ gates, constructs an MWIS-circuit of G with $\tau^{O(1)}$ gates in $\tau^{O(1)}$ time.*

In particular, a property analogous to automatizability of proof systems [6] holds for MWIS-circuits on bounded degree graphs and H-minor-free graphs.

---

[1] The stated result in [13] is that such a grid is a minor, but the same proof works directly to show that the constructed grid minor is also an induced minor. In particular, the proof in [13] does not use any edge deletions, and the corresponding result for bounded-genus graphs that it depends on [14] is already stated in terms of contraction to a graph that can be turned into a grid by removing vertices without decreasing treewidth by more than a constant factor.

MWIS-Formulas and Treedepth

We characterize optimal MWIS-formulas of bounded degree graphs.

▶ **Theorem 5.** *Let G be any graph with treedepth t and maximum degree d. Any MWIS-formula of G has $2^{\Omega(t/d)}$ gates.*

Again, Theorem 5 is optimal up to a factor d by the same construction as Theorem 1. As formulas can be thought of as bounded space analogies of circuits, Theorem 5 gives further evidence (in addition to e.g. [9, 26, 35]) supporting that while treewidth is the right parameter for CSP-like problems when equipped with unlimited space, treedepth is the right parameter when dealing with bounded space.

Obtaining a constant-factor single-exponential time parameterized approximation algorithm for treedepth is a well-known open problem [12], so while we know that the converse of Theorem 5 existentially holds in bounded degree graphs, we currently do not know how to construct such MWIS-formulas without having the treedepth decomposition as an input.

1.2 Techniques

Our main circuit complexity tool is an adaptation of a circuit decomposition lemma used in e.g. [22, 23, 37]. In particular, we show that this lemma can be adapted so that given an MWIS-circuit with τ gates of a graph with treewidth w it extracts a family of τ vertex separators each of size $\Omega(w)$. Once this family has been extracted, the main challenge for proving Theorem 1 is to show that if this family of separators is too small, there exists an independent set that intersects all of the separators. For this we use the lopsided Lovász Local Lemma [15], though we note that more elementary arguments would suffice to prove the theorem with a worse dependency on d. To extend the result from bounded degree graphs to H-minor-free graphs we use the minor model of the bounded degree induced minor with high treewidth to further control the structure of these separators.

For MWIS-formulas parameterized by treedepth t we similarly extract a family of 2τ vertex sets each of size $\Omega(t)$ from a τ-gate MWIS-formula, showing that if an independent set intersects all of these vertex sets the formula cannot compute it. The same application of the Local Lemma is used to prove that such an independent set indeed exists in low degree graphs if τ is too small. The argument for extracting the family from the formula is more ad-hoc than the argument for circuits.

1.3 Related Work

The convention of modeling dynamic programming algorithms as tropical circuits originates from the recent works of Jukna [24, 25], although some earlier results in monotone arithmetic circuit complexity apply also to tropical circuits [23]. In general, tropical circuit lower bounds imply lower bounds for monotone arithmetic circuits, but not necessarily the other way around [24]. In addition to the works of Jukna, the other works explicitly giving lower bounds for tropical circuits or formulas that we are aware of are [30, 31]. We are not aware of prior works on lower bounds for tropical circuits or formulas considering maximum weight independent set or the graph parameters treewidth or treedepth.

There are multiple worst-case hardness results related to different formulations of the independent set polynomial. In [7] it was shown that the multivariate independent set polynomial is VNP-complete. The univariate independent set polynomial is #P-hard to evaluate at every non-zero rational point [5], and more fine-grainedly its evaluation has $2^{\Omega(n/\log^3 n)}$ worst-case complexity assuming #ETH [21].

Chvátal has shown that a certain proof system for maximum independent set which naturally corresponds to branching algorithms requires exponential size proofs on almost all graphs that have the number of edges linear in the number of vertices [11].

Multiple worst-case lower bounds of form $n^{\Omega(w)}$ in limited models of computation for graph homomorphism problems of a pattern graph with treewidth w to a graph with n vertices are known [4, 26, 28]. In particular, recently it was shown that the worst-case monotone arithmetic circuit complexity of homomorphism polynomial is $\Theta(n^{w+1})$, and the worst-case monotone arithmetic formula complexity is $\Theta(n^t)$, where t is the treedepth of the pattern graph [26].

Recently, a lower bound of $2^{\Omega(w)}$ was shown for DNNF-compilation of monotone CNFs with primal treewidth w and bounded degree and arity, applying to all such CNFs [2]. We note that after the acceptance of this paper, we became aware of a reduction from MWIS-circuits to DNNFs that allows to prove a weaker version of our Theorem 1 via the result of [2]. In particular, the techniques of [2] yield an exponent of form $\Omega(w/2^d)$ instead of the best possible $\Omega(w/d)$ given in Theorem 1.

1.4   Organization

In Section 2 we present preliminaries on graph theory, define MWIS-circuits and prove simple lemmas on them, and discuss the lopsided Lovász Local Lemma and prove a lemma using it. In Section 3 we prove the lower bounds for MWIS-circuits parameterized by treewidth, i.e., Theorems 1 and 2. In Section 4 we prove the lower bound for MWIS-formulas parameterized by treedepth, i.e., Theorem 5. In Section 5 we give the construction that shows the optimality of Theorems 1 and 5 up to a factor d. We conclude and discuss future work in Section 6.

2   Preliminaries

2.1   Graphs

The vertex set of a graph G is denoted by $V(G)$ and the edge set by $E(G)$. The set of neighbors of a vertex v is denoted by $N(v)$ and the neighborhood of a vertex set X by $N(X) = \bigcup_{v \in X} N(v) \setminus X$. Closed neighborhoods are denoted by $N[v] = N(v) \cup \{v\}$ and $N[X] = N(X) \cup X$. The subgraph $G[X]$ induced by a vertex set $X \subseteq V(G)$ has $V(G[X]) = X$ and $E(G[X]) = \{\{u, v\} \in E(G) \mid u \in X \wedge v \in X\}$. We also use $G \setminus X = G[V(G) \setminus X]$ to denote induced subgraphs. An independent set of G is a vertex set I such that $G[I]$ has no edges. In particular, an empty set is an independent set.

A tree decomposition of a graph G is a tree T whose each vertex $i \in V(T)$ corresponds to a bag $B_i \subseteq V(G)$, satisfying that

1. $V(G) = \bigcup_{i \in V(T)} B_i$,

2. for each $\{u, v\} \in E(G)$ there is a bag B_i with $\{u, v\} \subseteq B_i$, and

3. for each $v \in V(G)$ the subtree of T induced by bags containing v is connected.

The width of a tree decomposition is $\max |B_i| - 1$ and the treewidth $tw(G)$ of a graph G is the minimum width over its tree decompositions.

A treedepth decomposition of a graph G is a rooted forest F with vertex set $V(F) = V(G)$, satisfying for each $\{u, v\} \in E(G)$ that u and v have an ancestor-descendant relation in F. The depth of F is the maximum number of vertices on a simple path from a root to a leaf in F. The treedepth $td(G)$ of a graph G is the minimum depth over its treedepth decompositions. Note that $tw(G) + 1 \leq td(G)$.

2.2 MWIS-Circuits

We start by giving a formal definition of a tropical circuit. Our definition is non-standard in that it does not allow any other input constants than 0, which we can w.l.o.g. assume in the context of maximum weight independent set. For a comprehensive treatment of tropical circuits and their relations to monotone Boolean and monotone arithmetic circuits see [24].

▶ **Definition 6.** A tropical circuit over variables X is a directed acyclic graph with in-degree of each vertex either 0 or 2. The vertices are called gates, the in-degree of a gate is called fan-in, and the out-degree of a gate is called fan-out. Each gate with fan-in 0 is labeled with a variable $x_i \in X$ or the constant 0 and each gate with fan-in 2 is labeled with either max or +. One gate is designated as the output gate. A tropical formula is a tropical circuit where each gate has fan-out at most 1.

With an assignment of real numbers to the variables X, a tropical circuit outputs a number computed by the output gate by natural semantics, i.e., a gate labeled with a variable x_i computes the value of x_i, a gate labeled with 0 computes 0, a gate labeled with + computes the sum of the values computed by its children, and a gate labeled with max computes the maximum of the values computed by its children. In particular, a tropical circuit computes a tropical polynomial in the variables X over the tropical $(\mathbb{R} \cup \{-\infty\}, \max, +)$ semiring. In the tropical semiring max corresponds to addition and + corresponds to multiplication, with $-\infty$ as the zero and 0 as the unit. We will refer to max as addition and to + as multiplication.

We define an *MWIS-polynomial* with the following simple lemma.

▶ **Lemma 7.** *Let G be a graph. A tropical circuit over variables $V(G)$ computes the weight of a maximum weight independent set of G for any assignment of real weights to the inputs if and only if for the tropical polynomial f computed by the circuit it holds that*
1. *each monomial of f is of form $v_1 \cdot \ldots \cdot v_l$, where $\{v_1, \ldots, v_l\}$ is an independent set of G and*
2. *for each independent set $\{v_1, \ldots, v_l\}$ of G there is a monomial $v_1 \cdot \ldots \cdot v_l$ in f, including the empty independent set corresponding to the empty product 0.*

Proof. For the if-direction, (1) guarantees that the value computed by the circuit is at most the weight of a maximum weight independent set and (2) guarantees that the value is at least the weight of a maximum weight independent set.

For the only if-direction, if some monomial would not be multilinear, i.e., include a factor v^2 for some vertex v, the output would be incorrect when assigning weight 1 to v and 0 to other vertices. If some monomial would be of form $v_1 \cdot \ldots \cdot v_l$, where $\{v_1, \ldots, v_l\}$ is not an independent set the output would be incorrect when assigning weight 1 to those v_i and 0 to others. Finally, if the output polynomial would not include $v_1 \cdot \ldots \cdot v_l$ as a monomial for some independent set $\{v_1, \ldots, v_l\}$ then the circuit would be incorrect when assigning weight 1 to vertices of this independent set and -1 to others.   ◀

An MWIS-polynomial of a graph G is a polynomial f satisfying (1) and (2) in Lemma 7. An MWIS-circuit of G is a tropical circuit that computes an MWIS-polynomial of G. An MWIS-formula of G is an MWIS-circuit of G that is a tropical formula.

We note that requiring the circuit to work for all real weights is not a strong assumption: Any MWIS-circuit that works for weights $\{0, 1\}$ can be turned into an MWIS-circuit that works for weights $\mathbb{R} \cup \{-\infty\}$ by replacing each input variable v_i by $\max(v_i, 0)$. In particular, the weight of an empty independent set is 0, so negative weights will never be used. Our assumption that the only constant available to the circuit is 0 can be justified by noting

that if an output monomial would contain a positive constant the circuit would be incorrect on the all-zero input, and that if an output monomial would contain a negative constant it should also occur without the constant. In particular, any other constants than 0 could be replaced by 0.

Next we make some simple observations on the structure of MWIS-circuits.

▶ **Definition 8.** A partial MWIS-polynomial is a polynomial f satisfying (1) in Lemma 7. A partial MWIS-circuit is a tropical circuit computing a partial MWIS-polynomial.

Note that by monotonicity of $(\max, +)$ computations we can assume that each gate of an MWIS-circuit computes a partial MWIS-polynomial and therefore each subcircuit is a partial MWIS-circuit.

▶ **Definition 9.** Let f be a partial MWIS-polynomial. We denote by $\mathrm{Sup}(f)$ the support of f, that is, the variables that occur in the monomials of f.

We also use $\mathrm{Sup}(g)$ for a gate g to denote the support of the polynomial computed by the gate. Note that each monomial of f corresponds to an independent set of $G[\mathrm{Sup}(f)]$.

The following property is the basis for proving lower bounds for MWIS-circuits.

▶ **Lemma 10.** *Let $f = g \cdot h$ be a partial MWIS-polynomial of a graph G. The sets $N[Sup(g)]$ and $Sup(h)$ are disjoint.*

Proof. If there was a vertex $v \in \mathrm{Sup}(g) \cap \mathrm{Sup}(h)$ then f would contain a monomial with a factor v^2. If there was a vertex $v \in \mathrm{Sup}(g)$ and $u \in \mathrm{Sup}(h)$ with $\{u, v\} \in E(G)$, then there would be a monomial in f containing a factor $u \cdot v$. ◀

We will say that a partial MWIS-polynomial f or a circuit computing f *computes* an independent set I if f contains the monomial $\prod_{v_i \in I} v_i$. In particular, an MWIS-polynomial computes every independent set.

2.3   Lopsided Lovász Local Lemma

The lopsided Lovász Local Lemma [15] (see [1] for the general version) is a method for showing that there is a non-zero probability that none of the events in a collection of events hold. In particular, we use it to show that independent sets satisfying certain requirements exist.

▶ **Definition 11.** Let $\mathcal{E}_1, \ldots, \mathcal{E}_n$ be events in a probability space. A graph Γ is a negative dependency graph of the events if its vertices are $V(\Gamma) = \{\mathcal{E}_1, \ldots, \mathcal{E}_n\}$ and for all events $\mathcal{E}_i$ and subsets $J \subseteq V(\Gamma) \setminus N(\mathcal{E}_i)$ it holds that $\Pr[\bigcup_{j \in J} \mathcal{E}_j \mid \mathcal{E}_i] \geq \Pr[\bigcup_{j \in J} \mathcal{E}_j]$.

In words, the negative dependency graph should capture all negative correlations between the events.

▶ **Proposition 12** ([1])**.** *Let $\mathcal{E}_1, \ldots, \mathcal{E}_n$ be a collection of events with a negative dependency graph Γ. If there exists real numbers $x_1, \ldots, x_n$ with $0 < x_i < 1$ such that for each i it holds that $\Pr[\mathcal{E}_i] \leq x_i \prod_{\mathcal{E}_j \in N(\mathcal{E}_i)} (1 - x_j)$, then $\Pr[\bigcap_{i=1}^{n} \overline{\mathcal{E}_i}] > 0$.*

2.4   Hitting Vertex Sets with Independent Sets

We prove a lemma which captures our use of the Local Lemma in Theorems 1 and 5. We spell out the constants to emphasize that they are not particularly high, although noting that a more careful proof could improve them a bit.

▶ **Lemma 13.** *Let G be a graph with maximum degree d and $\mathcal{F}$ a family of vertex subsets of G, each member of $\mathcal{F}$ containing at least k vertices. If $6|\mathcal{F}| \leq e^{k/(6d)}$, then there exists an independent set of G that intersects all sets in $\mathcal{F}$.*

Proof. We assume $d \geq 2$ as the lemma is easy to verify for $d \leq 1$. We use the Local Lemma to construct such an independent set. We let each vertex be in the independent set with probability $p = 1/(2d)$. Our bad events are $\mathcal{E}_e$ for each each edge e indicating that both endpoints of e are selected in the independent set, and $\mathcal{E}_A$ for each $A \in \mathcal{F}$ indicating that the independent set does not intersect A. The negative dependency graph is a bipartite graph connecting $\mathcal{E}_e$ to $\mathcal{E}_A$ if at least one of the endpoints of e is in A. In particular, note that the edge events $\mathcal{E}_e$ have non-negative correlation with each other and the vertex set events $\mathcal{E}_A$ also have non-negative correlation with each other. For all edge events $\mathcal{E}_e$ we choose $x_e = 1/(3d^2 + 1)$ and for all vertex set events $\mathcal{E}_A$ we choose $x_A = 1/(5|\mathcal{F}| + 1)$. Now, by Proposition 12, it suffices to verify that

$$\Pr[\mathcal{E}_e] = p^2 \leq x_e(1 - x_A)^{|\mathcal{F}|} \tag{1}$$

and

$$\Pr[\mathcal{E}_A] = (1 - p)^{|A|} \leq x_A(1 - x_e)^{|A|d} \tag{2}$$

hold whenever $6|\mathcal{F}| \leq e^{|A|/6d}$.

For (1), a lower bound for the right hand side is $e^{-1/5}/(3d^2 + 1)$, which can be verified to be greater than $p^2 = 1/(4d^2)$ when $d \geq 2$. For (2), an upper bound for the left hand side is $e^{-|A|/(2d)}$, and a lower bound for the right hand side is $x_A e^{-|A|d/(3d^2)}$, implying that (2) holds if $e^{-|A|/(2d)}e^{|A|/(3d)} \leq x_A$. This simplifies to $e^{-|A|/(6d)} \leq x_A \Leftrightarrow e^{|A|/(6d)} \geq 5|\mathcal{F}| + 1$. ◀

3 Treewidth and MWIS-Circuits

In this section we prove lower bounds for MWIS-circuits parameterized by treewidth, i.e., Theorems 1 and 2.

We use a witness of high treewidth due to Robertson-Seymour treewidth approximation algorithm [36]. A separation of a graph G is an ordered triple of vertex sets (A, S, B) such that A, S, B are disjoint, $A \cup S \cup B = V(G)$, and no vertex of A is adjacent to a vertex of B. The order of a separation (A, S, B) is $|S|$. A separation (A, S, B) is a balanced separation of a vertex set $X \subseteq V(G)$ if $|A \cap X| \leq 2|X|/3$ and $|B \cap X| \leq 2|X|/3$.

▶ **Lemma 14** ([36]). *If a graph G has treewidth at least $4k$, then there is a vertex set $X \subseteq V(G)$ such that any balanced separation of X in G has order at least k.*

The next lemma is our main tool to connect circuit complexity with treewidth. This lemma is an adaptation of a classical circuit decomposition lemma (e.g. Theorem 1 in [22], Lemma 3 in [37]). In our applications the vertex set X will be the set given by Lemma 14.

▶ **Lemma 15.** *Let G be a graph and $X \subseteq V(G)$ with $|X| \geq 2$. If there is an MWIS-circuit of G with τ gates, then we can write an MWIS-polynomial of G as $g_1 \cdot h_1 + \ldots + g_\tau \cdot h_\tau$, where for all i it holds that $|Sup(g_i) \cap X| \leq 2|X|/3$ and $|Sup(h_i) \cap X| \leq 2|X|/3$.*

Proof. Let $f + e$ be an MWIS-polynomial of G, where f can be computed by a tropical circuit with τ gates. (The term e is here for the induction argument. In the first step we can assume it to be empty.) We will show that there is an MWIS-polynomial $f' + g \cdot h + e$ of G, where f' can be computed by a tropical circuit with $\tau - 1$ gates, and $|Sup(g) \cap X| \leq 2|X|/3$ and $|Sup(h) \cap X| \leq 2|X|/3$. The lemma follows from this by induction.

If $|\text{Sup}(f) \cap X| \le 2|X|/3$ we are done. Otherwise, we traverse the circuit computing f down starting from the output gate, always choosing the one of the two child gates whose support has larger intersection with X, until we reach a gate v computing a polynomial f_v with $|X|/3 \le |\text{Sup}(f_v) \cap X| \le 2|X|/3$. Let $f_{v=-\infty}$ be the polynomial computed by the circuit when the value of the gate v is set to $-\infty$. Now we can write an MWIS-polynomial of G as $f_{v=-\infty} + f_v \cdot g + e$, for example by letting g be an MWIS-polynomial of $G \setminus N[\text{Sup}(f_v)]$. Now, we observe that $f_{v=-\infty}$ can be computed by a circuit with $\tau - 1$ gates. We also observe that the supports of f_v and g cannot intersect, and therefore $|\text{Sup}(g) \cap X| \le 2|X|/3$. ◄

3.1 Proof of Theorem 1

Now we complete the proof of Theorem 1 by putting Lemmas 13, 14, and 15 together.

▶ **Lemma 16.** *Let G be a graph with maximum degree d and treewidth at least $4k$. Any MWIS-circuit of G has at least $e^{k/(6d)}/6$ gates.*

Proof. Suppose there is an MWIS-circuit of G with τ gates. By Lemma 14 there is a vertex set $X \subseteq V(G)$ that does not admit a balanced separation of order less than k. By Lemma 15 we can write an MWIS-polynomial of G as $g_1 \cdot h_1 + \ldots + g_\tau \cdot h_\tau$, where for all i it holds that $|\text{Sup}(g_i) \cap X| \le 2|X|/3$ and $|\text{Sup}(h_i) \cap X| \le 2|X|/3$. Now, by Lemma 10 each multiplication $g_i \cdot h_i$ defines a balanced separation $(\text{Sup}(g_i), V(G) \setminus \text{Sup}(g_i \cdot h_i), \text{Sup}(h_i))$ of X. The order of such a separation is $|V(G) \setminus \text{Sup}(g_i \cdot h_i)|$, and therefore $|V(G) \setminus \text{Sup}(g_i \cdot h_i)| \ge k$. Note that $g_i \cdot h_i$ does not compute an independent set I if I intersects $V(G) \setminus \text{Sup}(g_i \cdot h_i)$. Therefore, by letting $\mathcal{F}$ be the collection of vertex sets $\{V(G) \setminus \text{Sup}(g_1 \cdot h_1), \ldots, V(G) \setminus \text{Sup}(g_\tau \cdot h_\tau)\}$, Lemma 13 shows that if $6\tau \le e^{k/(6d)}$ we can construct an independent set that is not computed by any of the multiplications, contradicting the assumption that we have an MWIS-circuit. ◄

3.2 Proof of Theorem 2

An *induced minor model* of a graph H in a graph G is a function $f : V(H) \to 2^{V(G)} \setminus \{\emptyset\}$, where $2^{V(G)}$ denotes the power set of $V(G)$, satisfying that
1. the sets $f(u)$ and $f(v)$ are disjoint for $u \ne v$,
2. for each $v \in V(H)$ the induced subgraph $G[f(v)]$ is connected, and
3. $\{u, v\} \in E(H)$ if and only if $N(f(u))$ intersects $f(v)$.

A graph G contains a graph H as an induced minor if and only if there is an induced minor model of H in G. For $v \in V(H)$ we call the induced subgraphs $G[f(v)]$ *clusters*.

First, we ensure that the maximum degree of each cluster is bounded.

▶ **Lemma 17.** *Let G be a graph that contains a graph H with maximum degree d as an induced minor. There is an induced minor model f of H in G such that the maximum degree of each cluster $G[f(v)]$ is at most d.*

Proof. Consider an induced minor model f of H in G and a cluster $G[f(v)]$ for some $v \in V(H)$. Because the degree of H is at most d, we can assign the cluster a set of at most d terminal vertices whose connectivity should be preserved in order to satisfy that f is an induced minor model of H in G. Now, we can remove from the cluster any vertices as long as the terminals stay connected. In particular, if there is a vertex u with degree $> d$ in $G[f(v)]$, then we can consider the shortest paths from u to the terminals, and remove from $G[f(v)]$ the vertices of $N(u) \cap G[f(v)]$ that do not participate in the shortest paths. This makes the degree of u in $G[f(v)]$ at most d. ◄

We also need the following lemma.

▶ **Lemma 18.** *Let I be an independent set selected uniformly at random from the set of all independent sets of a graph G with maximum degree d. For all $v \in V(G)$ it holds that $\Pr[v \in I] \geq 1/2^{d+1}$.*

Proof. For any set $J \subseteq N(v)$ it holds that $\Pr[I \cap N(v) = J] \leq \Pr[I \cap N(v) = \emptyset]$ because we can map any independent set I with $I \cap N(v) = J$ into an independent set $I \setminus N(v)$. Therefore $\Pr[I \cap N(v) = \emptyset] \geq 1/2^d$, so by observing that $\Pr[v \in I \mid I \cap N(v) = \emptyset] \geq 1/2$ we get $\Pr[v \in I] \geq 1/2^{d+1}$.   ◀

Next we finish the proof with similar arguments as in the proof of Theorem 1, but with a different kind of construction of the independent set with the Local Lemma. In this case the constants involved appear to be impractical.

▶ **Lemma 19.** *Let G be a graph that contains a graph H with maximum degree d and treewidth $4k$ as an induced minor. Any MWIS-circuit of G has $2^{\Omega(k/(d4^d))}$ gates.*

Proof. Let f be the induced minor model of H in G. First, by Lemma 17 we can assume that the maximum degree of each cluster $G[f(v)]$ is at most d. Now, by Lemma 14 we let X' be a vertex set of H that has no balanced separation of order less than k. Then we let X be a vertex set of G created by mapping each $v \in X'$ to an element of $f(v)$. For each balanced separation (A, S, B) of X in G, the set S must intersect at least k different clusters, because otherwise we could map it into a balanced separation of X' of order $< k$ in H. Therefore, by assuming that G has an MWIS-circuit with τ gates and applying Lemma 15 with the set X we write an MWIS-polynomial of G as $g_1 \cdot h_1 + \ldots + g_\tau \cdot h_\tau$, observing that for each i the set $S_i = V(G) \setminus \mathrm{Sup}(g_i \cdot h_i)$ intersects at least k different clusters. Now it remains to show that if τ is too small we can construct an independent set of G that intersects S_i for all i.

By removing vertices from each S_i we can assume that S_i contains only vertices in clusters, and moreover contains exactly one vertex from each cluster that it intersects. We use the Local Lemma to construct the independent set. First we select each cluster independently with probability $p = 1/(4d2^d)$, and then for each selected cluster $G[f(v)]$ we select an independent set uniformly at random from the set of all independent sets of $G[f(v)]$. By Lemma 18 each vertex of G that is in some cluster will appear in the independent set with probability at least $p/2^{d+1}$. Vertices in different clusters appear in it independently of each other.

Now our bad events are $\mathcal{E}_{\{u,v\}}$ for all $\{u,v\} \in E(H)$ indicating that both clusters $G[f(u)]$ and $G[f(v)]$ have been selected and $\mathcal{E}_i$ for each S_i indicating that the set S_i does not intersect the independent set. Our negative dependency graph has edges connecting each $\mathcal{E}_{\{u,v\}}$ to each $\mathcal{E}_i$ such that S_i intersects $f(u)$ or $f(v)$. It also has all edges between all events $\mathcal{E}_i$ because $\mathcal{E}_i$ and $\mathcal{E}_j$ can be negatively correlated if S_i and S_j intersect a common cluster.

For edges $\{u,v\} \in E(H)$ we let $x_{\{u,v\}} = 1/(15d^2 4^d + 1)$ and for sets S_i we choose $x_i = 1/(20\tau + 1)$. Now it suffices to verify that

$$\Pr[\mathcal{E}_{\{u,v\}}] = p^2 \leq x_{\{u,v\}}(1 - x_i)^\tau \tag{3}$$

and

$$\Pr[\mathcal{E}_i] \leq (1 - p/2^{d+1})^{|S_i|} \leq x_i(1 - x_i)^\tau (1 - x_{\{u,v\}})^{|S_i|d} \tag{4}$$

hold whenever $30\tau \leq e^{7|S_i|/(120d4^d)}$. We also assume that $d \geq 3$ since if $d \leq 2$ then the treewidth of H is at most 2.

For (3), a lower bound for the right hand side is $e^{-1/20}/(15d^2 4^d + 1)$, which is greater than $p^2 = 1/(16d^2 4^d)$ when $d \geq 2$. For (4), a lower bound for the right hand side is $x_i e^{-1/20} e^{-|S_i|d/(15d^2 4^d)}$ and an upper bound for the left hand side is $e^{-|S_i|/(8d4^d)}$, so it holds whenever $e^{-|S_i|/(8d4^d)} \leq x_i e^{-1/20} e^{-|S_i|d/(15d^2 4^d)}$ holds, which we can simplify to $e^{|S_i|(1/(15d4^d)-1/(8d4^d))} \leq x_i e^{-1/20}$, and finally to $e^{-7|S_i|/(120d4^d)} \leq 1/(20\tau + 1)e^{-1/20}$, which holds whenever $30\tau \leq e^{7|S_i|/(120d4^d)}$. ◀

4   Treedepth and MWIS-Formulas

For treedepth we are not aware of linear high-treedepth witnesses similar to what Lemma 14 is for treewidth. However, it turns out that we can use very basic properties of treedepth decompositions to establish the connection to formula complexity.

Recall that we denote the treedepth of a graph G with $td(G)$. The following properties follow from the definition of treedepth.

▶ **Proposition 20.** *Let G be a graph with treedepth $td(G)$. It holds that*
1. *$td(G \setminus \{v\}) \geq td(G) - 1$ for any $v \in V(G)$ and*
2. *$td(G)$ is the maximum of $td(G[C])$ over the connected components C of G.*

For our proof we need to introduce two definitions on MWIS-formulas. We start by defining *typical* independent sets of a partial MWIS-formulas.

▶ **Definition 21.** Let F be a partial MWIS-formula of a graph G. An independent set I of G is a typical independent set of F if for each multiplication gate g with $td(G[\mathrm{Sup}(g)]) \geq td(G)/2$ it holds that I intersects a connected component C of $G[\mathrm{Sup}(g)]$ with $td(G[C]) = td(G[\mathrm{Sup}(g)])$.

Note that by the property 2 of Proposition 20 such component indeed exists.

We also define the *separator* $\mathrm{Sep}(g)$ of a gate g. Note that an MWIS-formula forms a tree rooted at the output gate, so we will use standard tree terminology (parent, child, ancestor, descendant).

▶ **Definition 22.** The separator of the output gate o is $\mathrm{Sep}(o) = V(G) \setminus \mathrm{Sup}(o)$. The separator of a gate g whose parent p is a multiplication gate is $\mathrm{Sep}(g) = \mathrm{Sep}(p)$. The separator of a gate g whose parent p is a sum gate is $\mathrm{Sep}(g) = \mathrm{Sep}(p) \cup \mathrm{Sup}(p) \setminus \mathrm{Sup}(g)$.

With the definitions of typical independent sets and separators of gates, we can state the following lemma which will be applied with Lemma 13 to prove our lower bound.

▶ **Lemma 23.** *Let G be a graph with $td(G) \geq 2$ and F a partial MWIS-formula of G. If I is a typical independent set of F and intersects $\mathrm{Sep}(g)$ for each gate g with $|\mathrm{Sep}(g)| \geq td(G)/2$, then F does not compute I.*

Proof. Let F be such a formula and I such an independent set. We say that a gate g of F is redundant if F computes I if and only if F without g computes I. First, note that all gates g such that I intersects $\mathrm{Sep}(g)$ are redundant because by the definition of separator there is an ancestor gate g' of g with a sum gate parent p such that none of the monomials M contributed from g' to p have $M = \prod_{v_i \in I \cap \mathrm{Sup}(p)} v_i$, implying that g' is redundant and thus all of its descendants are redundant.

Now, we prove by induction starting from the leaves that every gate g of F for which $|\mathrm{Sep}(g)| + td(G[\mathrm{Sup}(g)]) \geq td(G)$ holds is redundant. First, for all such gates g with $td(G[\mathrm{sup}(g)]) \leq td(G)/2$, including all leaves, we have that $|\mathrm{Sep}(g)| \geq td(G)/2$, making g redundant by our definition of I. For a sum gate g and its child c we have by property 1

of Proposition 20 that $td(G[\mathrm{Sup}(c)]) \geq td(G[\mathrm{Sup}(g)]) - |\mathrm{Sup}(g) \setminus \mathrm{Sup}(c)|$, rearranging to $td(G[\mathrm{Sup}(c)] \geq td(G[\mathrm{Sup}(g)]) - |\mathrm{Sep}(c)| + |\mathrm{Sep}(g)|$, and finally to $td(G[\mathrm{Sup}(c)]) + |\mathrm{Sep}(c)| \geq td(G[\mathrm{Sup}(g)]) + |\mathrm{Sep}(g)|$. This implies that if $|\mathrm{Sep}(g)| + td(G[\mathrm{Sup}(g)]) \geq td(G)$ then both children of g are redundant, making g redundant. For a multiplication gate g with $td(G[\mathrm{Sup}(g)]) \geq td(G)/2$ it follows from the typicality assumption that there is a child c of g with $td(G[\mathrm{Sup}(c)]) + |\mathrm{Sep}(c)| = td(G[\mathrm{Sup}(g)]) + |\mathrm{Sep}(g)|$ such that g is redundant if c is redundant. Therefore the induction works, and because $|\mathrm{Sep}(o)| + td(G[\mathrm{Sup}(o)]) \geq td(G)$ holds for the output gate o, the output gate is redundant and therefore the formula does not compute I. ◀

Now the only thing left to complete the proof of Theorem 5 is to show that if a formula has less than $2^{\Omega(td(G)/d)}$ gates then we can construct an independent set that is typical for the formula and intersects $\mathrm{Sep}(g)$ whenever $|\mathrm{Sep}(g)| \geq td(G)/2$. For an independent set to be typical it suffices that it intersects $\mathrm{Sup}(g)$ for all gates g with $|\mathrm{Sup}(g)| \geq td(G)/2$. Therefore it suffices to apply Lemma 13 with $\mathcal{F}$ consisting of $\mathrm{Sep}(g)$ for all $|\mathrm{Sep}(g)| \geq td(G)/2$ and $\mathrm{Sup}(g)$ for all $|\mathrm{Sup}(g)| \geq td(G)/2$. This yields a lower bound of $e^{td(G)/(12d)}/12$ for the number of gates.

5   Optimality of Theorems 1 and 5

We show that for each pair w, d we can construct a graph with treewidth $\Omega(w)$ and maximum degree $O(d)$ that admits an MWIS-formula with $d2^{w/d}$ gates.

If $d > w$ then a d-clique does the job. Otherwise, we take a bounded degree expander with w/d vertices, having treewidth $\Omega(w/d)$, constructible by e.g. [19]. We replace each vertex of the expander with a d-clique (which will be referred to as cluster) such that each vertex of a cluster is connected to each vertex of the clusters of the adjacent vertices. We denote the constructed graph with $G_{w,d}$

▶ **Proposition 24.** *The graph $G_{w,d}$ has treewidth $\Omega(w)$, maximum degree $O(d)$, and admits an MWIS-formula with $d2^{w/d}$ gates.*

Proof. The maximum degree is at most $(d + 1)$ times the maximum degree of the original bounded degree expander. The treewidth is $\Omega(w)$ because if a balanced separator contains one vertex from a cluster it must contain all vertices of the cluster.

Note that by a simple recursion any n-vertex graph admits an MWIS-formula with at most 2^n gates, so the original expander admits an MWIS-formula with $2^{w/d}$ gates. We can construct an MWIS-formula of $G_{w,d}$ by taking the MWIS-formula of the original expander and replacing each leaf corresponding to a vertex v with a d-gate construction computing the maximum over the vertices of the cluster of v. ◀

6   Conclusions and Future Work

We investigated the tropical circuit complexity of maximum weight independent set. Our initial motivation for this was the fact that lower bounds for tropical circuits imply lower bounds for many actual algorithmic techniques for maximum weight independent set that are widely used in both theory and practice. We showed that in bounded degree graphs optimal MWIS-circuits are characterized by treewidth and optimal MWIS-formulas are characterized by treedepth. We generalized the result for MWIS-circuits to apply beyond bounded degree

graphs, to a graph class that includes all planar graphs, and more generally all H-minor-free graphs. The constants hidden by the Ω-notation in Theorems 1 and 5 are somewhat practical even though we did not specifically optimize them. For example, Theorem 1 shows that any MWIS-circuit of the 5000×5000-grid has at least 10^{21} gates.

We identify some technical barriers for extending the results. First, we note that Lemma 13 is not effective in graphs with maximum degree higher than k: If $|N(v)| \geq k$, we can add $N(v)$ to $\mathcal{F}$ to force the independent set to avoid v, essentially forcing us to work with $G \setminus \{v\}$. Indeed an example of a graph with high treewidth and no small MWIS-circuits for which Lemma 13 is unsuitable is a clique with each edge subdivided. In some cases, including H-minor-free graphs and the subdivided clique, this barrier can be circumvented with Theorem 2 by using a bounded degree induced minor with high treewidth. We also note that our proofs do not exploit the fact that the separators given by Lemma 15 are balanced beyond just the size bound.

The subdivided clique does not exclude any fixed graph as a minor, so the fact that Theorem 2 works also for proving a lower bound for it seems to indicate that Theorem 2 is more powerful than what is captured by Corollary 3. We are in fact not aware of graph families for which a $2^{\Omega(w)}$ lower bound can be proved but Theorem 2 does not apply.

An interesting general direction for future work could be to prove Corollary 4 for as large graph classes as possible, starting by extending the $2^{\Omega(w)}$ lower bound as far as possible. In particular, H-topological-minor-free graphs generalize both bounded degree and H-minor-free graphs [33], so proving a $2^{\Omega(w)}$ lower bound for them seems like a natural next step. Even more generally, it could be that such a lower bound could even apply to all bounded degeneracy graphs. We hope that this line of work will lead to new insights on the structure of independent sets that could even be useful for positive results on algorithms for maximum weight independent set.

References

1   Noga Alon and Joel H. Spencer. *The Probabilistic Method, Third Edition.* Wiley-Interscience series in discrete mathematics and optimization. Wiley, 2008.

2   Antoine Amarilli, Florent Capelli, Mikaël Monet, and Pierre Senellart. Connecting knowledge compilation classes and width parameters. *Theory of Computing Systems*, 64(5):861–914, 2020. `doi:10.1007/s00224-019-09930-2`.

3   Stefan Arnborg and Andrzej Proskurowski. Linear time algorithms for NP-hard problems restricted to partial k-trees. *Discrete Applied Mathematics*, 23(1):11–24, 1989. `doi:10.1016/0166-218X(89)90031-0`.

4   Per Austrin, Petteri Kaski, and Kaie Kubjas. Tensor network complexity of multilinear maps. In Avrim Blum, editor, *Proceedings of the 10th Innovations in Theoretical Computer Science Conference*, volume 124 of *LIPIcs*, pages 7:1–7:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. `doi:10.4230/LIPIcs.ITCS.2019.7`.

5   Markus Bläser and Christian Hoffmann. On the complexity of the interlace polynomial. In Susanne Albers and Pascal Weil, editors, *Proceedings of the 25th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1 of *LIPIcs*, pages 97–108. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2008. `doi:10.4230/LIPIcs.STACS.2008.1337`.

6   Maria Luisa Bonet, Toniann Pitassi, and Ran Raz. On interpolation and automatization for Frege systems. *SIAM Journal on Computing*, 29(6):1939–1967, 2000. `doi:10.1137/S0097539798353230`.

7   Irénée Briquel and Pascal Koiran. A dichotomy theorem for polynomial evaluation. In Rastislav Královic and Damian Niwinski, editors, *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science*, volume 5734 of *Lecture Notes in Computer Science*, pages 187–198. Springer, 2009. `doi:10.1007/978-3-642-03816-7_17`.

8   Binh-Minh Bui-Xuan, Jan Arne Telle, and Martin Vatshelle. Boolean-width of graphs. *Theoretical Computer Science*, 412(39):5187–5204, 2011. `doi:10.1016/j.tcs.2011.05.022`.

9   Li-Hsuan Chen, Felix Reidl, Peter Rossmanith, and Fernando Sánchez Villaamil. Width, depth, and space: Tradeoffs between branching and dynamic programming. *Algorithms*, 11(7):98, 2018. `doi:10.3390/a11070098`.

10   Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. On the maximum weight independent set problem in graphs without induced cycles of length at least five. *SIAM Journal on Discrete Mathematics*, 34(2):1472–1483, 2020. `doi:10.1137/19M1249473`.

11   V. Chvátal. Determining the stability number of a graph. *SIAM Journal on Computing*, 6(4):643–662, 1977. `doi:10.1137/0206046`.

12   Wojciech Czerwinski, Wojciech Nadara, and Marcin Pilipczuk. Improved bounds for the excluded-minor approximation of treedepth. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *Proceedings of the 27th Annual European Symposium on Algorithms*, volume 144 of *LIPIcs*, pages 34:1–34:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. `doi:10.4230/LIPIcs.ESA.2019.34`.

13   Erik D. Demaine and MohammadTaghi Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica*, 28(1):19–36, 2008. `doi:10.1007/s00493-008-2140-4`.

14   Erik D. Demaine, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos. The bidimensional theory of bounded-genus graphs. *SIAM J. Discret. Math.*, 20(2):357–371, 2006. `doi:10.1137/040616929`.

15   Paul Erdős and Joel Spencer. Lopsided Lovász local lemma and latin transversals. *Discrete Applied Mathematics*, 30(2-3):151–154, 1991. `doi:10.1016/0166-218X(91)90040-4`.

16   Fedor V. Fomin and Petr A. Golovach. Subexponential parameterized algorithms and kernelization on almost chordal graphs. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *Proceedings of the 28th Annual European Symposium on Algorithms*, volume 173 of *LIPIcs*, pages 49:1–49:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.ESA.2020.49`.

17   Fedor V. Fomin, Ioan Todinca, and Yngve Villanger. Large induced subgraphs via triangulations and CMSO. *SIAM Journal on Computing*, 44(1):54–87, 2015. `doi:10.1137/140964801`.

18   Eugene C. Freuder and Michael J. Quinn. Taking advantage of stable sets of variables in constraint satisfaction problems. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 1076–1078. Morgan Kaufmann, 1985. URL: `http://ijcai.org/Proceedings/85-2/Papers/082.pdf`.

19   Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981. `doi:10.1016/0022-0000(81)90040-4`.

20   Peter Gartland and Daniel Lokshtanov. Independent set on P_k-free graphs in quasi-polynomial time. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science*, pages 613–624. IEEE, 2020. `doi:10.1109/FOCS46700.2020.00063`.

21   Christian Hoffmann. Exponential time complexity of weighted counting of independent sets. In Venkatesh Raman and Saket Saurabh, editors, *Proceedings of the 5th International Symposium on Parameterized and Exact Computation*, volume 6478 of *Lecture Notes in Computer Science*, pages 180–191. Springer, 2010. `doi:10.1007/978-3-642-17493-3_18`.

22   Laurent Hyafil. On the parallel evaluation of multivariate polynomials. *SIAM Journal on Computing*, 8(2):120–123, 1979. `doi:10.1137/0208010`.

23   Mark Jerrum and Marc Snir. Some exact complexity results for straight-line computations over semirings. *Journal of the ACM*, 29(3):874–897, 1982. `doi:10.1145/322326.322341`.

24   Stasys Jukna. Lower bounds for tropical circuits and dynamic programs. *Theory of Computing Systems*, 57(1):160–194, 2015. `doi:10.1007/s00224-014-9574-4`.

25   Stasys Jukna. Tropical complexity, sidon sets, and dynamic programming. *SIAM Journal on Discrete Mathematics*, 30(4):2064–2085, 2016. `doi:10.1137/16M1064738`.

26 Balagopal Komarath, Anurag Pandey, and C. S. Rahul. Graph homomorphism polynomials: Algorithms and complexity. *CoRR*, abs/2011.04778, 2020. `arXiv:2011.04778`.

27 Casimir Kuratowski. Sur le probleme des courbes gauches en topologie. *Fundamenta mathematicae*, 15(1):271–283, 1930.

28 Yuan Li, Alexander A. Razborov, and Benjamin Rossman. On the AC^0 complexity of subgraph isomorphism. *SIAM Journal on Computing*, 46(3):936–971, 2017. `doi:10.1137/14099721X`.

29 Daniel Lokshantov, Martin Vatshelle, and Yngve Villanger. Independent set in P_5-free graphs in polynomial time. In Chandra Chekuri, editor, *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 570–581. SIAM, 2014. `doi:10.1137/1.9781611973402.43`.

30 Meena Mahajan, Prajakta Nimbhorkar, and Anuj Tawari. Computing the maximum using (min, +) formulas. In Kim G. Larsen, Hans L. Bodlaender, and Jean-François Raskin, editors, *Proceedings of the 42nd International Symposium on Mathematical Foundations of Computer Science*, volume 83 of *LIPIcs*, pages 74:1–74:11. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. `doi:10.4230/LIPIcs.MFCS.2017.74`.

31 Meena Mahajan, Prajakta Nimbhorkar, and Anuj Tawari. Shortest path length with bounded-alternation (min, +) formulas. *International Journal of Advances in Engineering Sciences and Applied Mathematics*, 11(1):68–74, 2019. `doi:10.1007/s12572-018-0229-6`.

32 J. W. Moon and L. Moser. On cliques in graphs. *Israel Journal of Mathematics*, 3(1):23–28, 1965.

33 Jaroslav Nesetril and Patrice Ossona de Mendez. On nowhere dense graphs. *European Journal of Combinatorics*, 32(4):600–617, 2011. `doi:10.1016/j.ejc.2011.01.006`.

34 Patric R. J. Östergård. A fast algorithm for the maximum clique problem. *Discrete Applied Mathematics*, 120(1-3):197–207, 2002. `doi:10.1016/S0166-218X(01)00290-6`.

35 Michał Pilipczuk and Marcin Wrochna. On space efficiency of algorithms working on structural decompositions of graphs. *ACM Transactions on Computation Theory*, 9(4):18:1–18:36, 2018. `doi:10.1145/3154856`.

36 Neil Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986. `doi:10.1016/0196-6774(86)90023-4`.

37 L. G. Valiant. Negation can be exponentially powerful. *Theoretical Computer Science*, 12:303–314, 1980. `doi:10.1016/0304-3975(80)90060-2`.