

On Tseitin formulas, read-once branching programs and treewidth

Ludmila Glinskikh and Dmitry Itsykson

lglinskikh@gmail.com, dmitrits@pdmi.ras.ru

St. Petersburg Department of V.A. Steklov Institute of Mathematics of the Russian Academy of Sciences,
Fontanka 27, St. Petersburg, Russia

March 25, 2019

Abstract

We show that any nondeterministic read-once branching program that decides a satisfiable Tseitin formula based on an $n \times n$ grid graph has size at least $2^{\Omega(n)}$. Then using the Excluded Grid Theorem by Robertson and Seymour we show that for arbitrary graph $G(V, E)$ any nondeterministic read-once branching program that computes a satisfiable Tseitin formula based on G has size at least $2^{\Omega(\text{tw}(G)^\delta)}$ for all $\delta < 1/36$, where $\text{tw}(G)$ is the treewidth of G (for planar graphs and some other classes of graphs the statement holds for $\delta = 1$). We also show an upper bound of $O(|E|2^{\text{pw}(G)})$, where $\text{pw}(G)$ is the pathwidth of G .

We apply the mentioned results to the analysis of the complexity of derivations in the proof system OBDD(\wedge , reordering) and show that any OBDD(\wedge , reordering)-refutation of an unsatisfiable Tseitin formula based on a graph G has size at least $2^{\Omega(\text{tw}(G)^\delta)}$.

1 Introduction

This paper continues the study of representation of satisfiable Tseitin formulas by read-once branching programs.

A Tseitin formula $\text{TS}_{G,c}$ [19] is defined for every undirected graph $G(V, E)$ and labelling function $c : V \rightarrow \{0, 1\}$. We introduce a propositional variable for every edge of G . The Tseitin formula $\text{TS}_{G,c}$ represents a linear system over the field $\text{GF}(2)$ that for every vertex $v \in V$ states that the sum of all edges adjacent to v equals $c(v)$. It is well known that a Tseitin formula is satisfiable if and only if the sum of values of the labelling function for all vertices in every connected component is even [20].

In 2017 Itsykson et al. [12] showed that any OBDD representing satisfiable Tseitin formulas based on d -regular expanders on n vertices has size at least $2^{\Omega(n)}$. Then Glinskikh and Itsykson [9] extended this lower bound to nondeterministic read-once branching programs (1-NBP).

In this paper we consider an $n \times n$ grid and study the complexity of representation of Tseitin formulas based on it by read-once branching programs. In Theorem 3.1 we prove that any 1-NBP computing a satisfiable Tseitin formula based on an $n \times n$ grid has size $2^{\Omega(n)}$. Although an $n \times n$ grid graph has some edge-expansion properties, we could not prove the lower bound based only on these properties; our proof requires careful analysis and we use the geometric properties of the grid.

As an important corollary we establish a connection between the complexity of 1-NBP representation of a satisfiable Tseitin formula and the treewidth of the underlying graph. The treewidth is one of the most important structural measures of a graph and it is one of the main parametrizations for computational graph problems. Theorem 4.1 states that any 1-NBP computing a satisfiable Tseitin formula $\text{TS}_{G,c}$ has size at least $2^{\Omega(\text{tw}(G)^\delta)}$ for all $\delta < 1/36$, where $\text{tw}(G)$ denotes the treewidth of the graph G . The proof is based on the Excluded Grid Theorem by Robertson and Seymour [18]: there is a function g such that if a graph G has treewidth at least $g(t)$, then G contains a grid of size $t \times t$ as a minor. Recent results of Chekuri and

Chuzhoy [4] and [5] give polynomial upper bound on the function g . Hence we know that every graph G has a $t \times t$ grid as a minor, where $t = \Omega(\text{tw}(G)^\delta)$ for all $\delta < 1/36$. For several classes of graphs it is possible to improve the value of δ , for example, for planar graphs $\delta = 1$ [10], [17]. Thus Theorem 4.1 is followed by Lemma 4.3 stating that if H is a minor of G , then for every S and for every 1-NBP of size S that computes a satisfiable Tseitin formula $\text{TS}_{G,c}$ there is an 1-NBP that computes a satisfiable Tseitin formula $\text{TS}_{H,c'}$ of size at most S . This lemma is proved separately for every operation: an edge deletion, a vertex deletion and an edge contraction. We use the non-determinism in the case of an edge contraction: we replace nodes labelled with contracted edges by guessing nodes.

In Theorem 4.5 we show that for every satisfiable Tseitin formula based on a graph $G(V, E)$ has an OBDD of size $O(|E|2^{\text{pw}(G)})$, where $\text{pw}(G)$ is the pathwidth of G (note that the pathwidth differs from the treewidth by at most a logarithmic factor: $\text{tw}(G) \leq \text{pw}(G) \leq O(\text{tw}(G) \log |V|)$). Since the pathwidth of an $n \times n$ grid is $O(n)$, our upper and lower bounds for grids match up to a constant in the exponent.

There are several known approaches to defining the treewidth of CNF formulas. Ferrara, Pan and Vardi [7] considered a graph on variables of a CNF formula where two variables are connected iff they share a common clause. They proved that if a graph associated with CNF formula has the treewidth t , then the formula has an OBDD of size $n^{O(t)}$ (it is very similar to Theorem 4.5 but it uses another notion of treewidth). Razgon [16] showed that this bound is tight and there is a family of CNF formulas with the treewidth at most k that requires 1-NBP of size $n^{\Omega(k)}$. In the case of a Tseitin formula $\text{TS}_{G,c}$, the associated graph is the edge-graph of G , where the vertices are the edges of G and two edges are connected iff they are incident to the same vertex of G . For example, if G is a star on $n + 1$ vertices (a star is a tree and, hence, it has the treewidth 1), then the edge-graph is the complete graph K_{n-1} and it has the treewidth $n - 2$.

Applications to proof complexity. The interest of the study of Tseitin formulas comes from the propositional proof complexity; unsatisfiable Tseitin formulas are one of the basic examples of hard formulas for many proof systems.

The study of representations of satisfiable Tseitin formulas by read-once branching program was motivated by the study of proof systems based on OBDDs introduced by Atserias, Kolaitis and Vardi [2]. Itsykson et al. [12] studied the proof system $\text{OBDD}(\wedge, \text{reordering})$; a proof of unsatisfiability of a CNF formula φ in this proof system is a sequence of OBDDs: D_1, D_2, \dots, D_s such that D_s is a constant false OBDD and for all $i \in [s]$, D_i either represents a clause of φ , or represents the conjunction of $D_k \wedge D_\ell$, where $k, \ell < i$ and D_i, D_k and D_ℓ use the same order, or represents the same function as D_ℓ but using another order, where $\ell < i$. The paper [12] gives an exponential lower bound on size of $\text{OBDD}(\wedge, \text{reordering})$ -refutations of unsatisfiable Tseitin formulas based on constant degree expanders. The lower bound proof is organized as follows: for any refutation of Tseitin formula $\text{TS}_{G,c}$ of size S it is possible to construct an OBDD of size at most S^2 representing a satisfiable Tseitin formula $\text{TS}_{G',c'}$, where G' is a graph obtained from G by the deletion of several edges. Thus it is sufficient to prove lower bound on the size of OBDD representation of $\text{TS}_{G',c'}$. We adapt this approach and show in Theorem 5.3 that our results imply that any $\text{OBDD}(\wedge, \text{reordering})$ -refutation of an unsatisfiable Tseitin formula $\text{TS}_{G,c}$ has size at least $2^{\Omega(\text{tw}(G)^\delta)}$, where δ is a constant as above. In particular we get a lower bound $2^{\Omega(n)}$ on the complexity of $\text{OBDD}(\wedge, \text{reordering})$ -refutations of Tseitin formulas based on the $n \times n$ grid.

The recent paper by Buss et al. [3] shows that this proof system cannot be polynomially simulated by Resolution and even by Cutting Planes. The paper shows that any Resolution proof of Tseitin formula based on the complete graph on $\log n$ vertices $K_{\log n}$ has size at least $2^{\Omega(\log^2 n)}$, while it has an $\text{OBDD}(\wedge, \text{reordering})$ -refutation of polynomial size. It is well known that the size of the shortest regular Resolution proof of any unsatisfiable CNF formula ϕ equals the size of the minimal read-once branching program for the following search problem Search_ϕ : given an assignment of variables of ϕ , find a clause that is refuted by this assignment [15], [14]. Our upper bound implies that satisfiable $\text{TS}_{K_{\log n}, c}$ can be computed by an OBDD of size $\text{poly}(n)$. Thus we have that computing of $\text{Search}_{\text{TS}_{K_{\log n}, c}}$ for an unsatisfiable $\text{TS}_{K_{\log n}, c}$ is superpolynomially harder than computing of a satisfiable $\text{TS}_{K_{\log n}, c'}$ for read-once branching programs.

Tseitin formulas based on the grid graphs were studied in proof complexity. The first superpolynomial lower bound for regular resolution was proved for grid graphs in 1968 by Tseitin [19]. In 1987 Urquhart

proved a lower bound for Tseitin formulas based on expanders in unrestricted Resolution [20] but tight lower bounds for grids were proved by Dantchev and Riis only in 2001 [6]. In the recent paper [11] Hastad proved lower bound on Bounded depth Frege refutations for Tseitin formulas based on $n \times n$ grid graphs that implies that polynomial size Frege proofs of such formulas should use formulas with almost logarithmic depth.

The treewidth was also studied in the context of resolution refutations of Tseitin formulas. Alekhovich and Razborov [1] considered a hypergraph that corresponds to every CNF formula, where variables are vertices and clauses as sets of variables form hyperedges. For Tseitin formulas the branch-width of this hypergraph is up to a constant factor equal to the Resolution width [1]. For constant degree graphs the treewidth is equal to the branch-width of the hypergraph up to a multiplicative constant. Galesi, Talebanfard and Torán in the recent paper [8] consider cop-robber games on graphs that is very similar to games characterising the treewidth, they used such games in an analysis of the complexity parameters of resolution refutations of Tseitin formulas.

Further research. It is interesting to discover whether there are examples of graphs such that the Resolution complexity of Tseitin formulas based on them is superpolynomially smaller than the shortest size of 1-BP of corresponding satisfiable Tseitin formulas.

2 Preliminaries

Branching programs. A deterministic branching program (BP) is a form of representation of Boolean functions. A Boolean function $f(x_1, x_2, \dots, x_n)$ is represented by a directed acyclic graph with exactly one source and two sinks. All nodes except sinks are labeled with a variable; every internal node has exactly two outgoing edges: one is labeled with 1 and the other is labeled with 0. One of the sinks is labeled with 1 and the other is labeled with 0. The value of the function for given values of variables is evaluated as follows: we start a path from the source such that for every node on its path we go along the edge that is labeled with the value of the corresponding variable. This path will end in a sink. The label of this sink is the value of the function.

A nondeterministic branching program (NBP) differs from a deterministic in the way that we also allow guessing nodes that are unlabeled and have two outgoing unlabeled edges. So nondeterministic branching program may have three types of nodes: guessing nodes, nodes labeled with a variable (we call them just labeled nodes) and two sinks; the source is either a guessing node or a labeled node. The result of a function represented by a nondeterministic branching program for given values of variables equals 1, if there exists at least one path from the source to the sink labeled with 1 such that for every node labeled with a variable on its path we go along an edge that is labeled with the value of the corresponding variable (for guessing nodes we are allowed to choose any of two outgoing edges). Note that deterministic branching programs constitute a special case of nondeterministic branching programs.

A deterministic or nondeterministic branching program is (syntactic) read- k (k -BP or k -NBP) if every path from the source to a sink contains at most k occurrences of each variable.

Let π be a permutation of the set $\{1, \dots, n\}$ (an order). A π -ordered binary decision diagram is a 1-BP such that on every path from the source to a sink variable $x_{\pi(i)}$ can not appear before $x_{\pi(j)}$ if $i > j$. An ordered binary decision diagram (OBDD) is a π -ordered binary decision diagram for some π .

Lemma 2.1 ([21]). *Assume that Boolean functions f_1 and f_2 have π -ordered OBDDs of sizes k_1 and k_2 respectively, then 1) $f_1 \wedge f_2$ has a π -ordered OBDD of size at most $k_1 k_2$; 2) for any partial substitution ρ , $f_1|_\rho$ has a π -ordered OBDD of size at most k_1 .*

Tseitin formulas. Let $G(V, E)$ be an undirected graph without loops but possibly with multiple edges, $c : V \rightarrow \{0, 1\}$ be a labeling function that matches every vertex with a Boolean value. We associate every edge $e \in E$ with a propositional variable x_e . A Tseitin formula $TS_{G,c}$ based on a graph G and a labeling

function c is the conjunction of the following conditions: for every vertex v the sum of variables x_e for all edges e that are incident to v equals $c(v)$ modulo 2. More formally: $\bigwedge_{v \in V} \left(\sum_{e \text{ is incident to } v} x_e = c(v) \pmod{2} \right)$.

Usually, Tseitin formulas are written in the CNF. If the maximal degree of a graph is upper bounded by a constant d , then a sum modulo 2 can be written as a d -CNF of size at most 2^d , hence the size of CNF representation of $\text{TS}_{G,c}$ is at most $O(2^d n)$.

We will use the following criterion of the satisfiability of Tseitin formulas:

Lemma 2.2 ([20]). *A Tseitin formula $\text{TS}_{G,c}$ is satisfiable if and only if for every connected component of the graph G the sum of values of the function c for all of the vertices is even. I.e., for every connected component U the following holds: $\sum_{v \in U} c(v) = 0 \pmod{2}$.*

Remark 2.3. *Note that a substitution of a value to a variable $x_e := \alpha$ transforms Tseitin formula $\text{TS}_{G,c}$ to a Tseitin formula $\text{TS}_{G',c'}$, where graph G' is obtained from the graph G by deleting the edge e , c' equals c in every vertex except two vertices that are incident to edge e . On these two vertices the values of c and c' differ by α .*

For a graph $G(V, E)$ let $k_G(l)$ be the maximal number of connected components that can be obtained from G by the deletion of l edges. The following lower bound on the size of 1-NBP for satisfiable Tseitin formula is known:

Lemma 2.4 ([9], Corollary 20). *For every connected graph $G(V, E)$ and arbitrary $1 \leq l \leq |E|$ any 1-NBP evaluating a satisfiable Tseitin formula $\text{TS}_{G,c}$ has size at least $2^{|V| - k_G(l) - k_G(|E| - l) + 1}$.*

Proof sketch. If a graph $H(U, F)$ consists of k connected components, then a satisfiable Tseitin formula $\text{TS}_{H,f}$ has exactly $2^{|F| - |U| + k}$ satisfying assignments ([9], Lemma 2). For every l we estimate the number of nodes of an 1-NBP for $\text{TS}_{G,c}$ on the level l . The graph G is connected, hence $\text{TS}_{G,c}$ has exactly $2^{|E| - |V| + 1}$ satisfying assignments. For every satisfying assignment of $\text{TS}_{G,c}$ we consider an accepting path of the 1-NBP corresponding to it. We consider the beginnings of these paths of length l . The number of accepting paths with the same beginning of length l is at most $2^{|E| - l - |V| + k_G(|E| - l)}$ (it is an upper bound on the number of satisfying assignments for Tseitin formulas on subgraph of G with $|E| - l$ edges). Thus, the number of different beginnings of length l of the accepting paths is at least $2^{l - k_G(|E| - l) + 1}$. The number of different beginnings of length l of accepting paths that go through a fixed vertex on the level l is at most $2^{l - |V| + k_G(|E| - l)}$ (it is an upper bound on the number of satisfying assignments for Tseitin formulas on subgraph of G with l edges). Finally, the number of vertices on the l -th level is at least $2^{|V| - k_G(l) - k_G(|E| - l) + 1}$. \square

Lemma 2.5. *Let G be an undirected graph, c and c' be such that Tseitin formulas $\text{TS}_{G,c}$ and $\text{TS}_{G,c'}$ are satisfiable. Then sizes of minimum-size 1-NBPs (1-BPs and OBDDs) for $\text{TS}_{G,c}$ and $\text{TS}_{G,c'}$ are equal.*

Proof. We will show that the first Tseitin formula can be obtained from the second by the changing some variables by their negations and, hence a branching program for the first Tseitin formula can be obtained from a branching program for the second Tsetin formula just by changing labels of outgoing edges for some nodes. Consider one connected component U of G . By Lemma 2.2 for the both Tseitin formulas the sum of labels that correspond to this component is even: $\sum_{v \in U} c(v) = \sum_{v \in U} c'(v) = 0$. Consider a set of vertices $U' = \{v \in U \mid c(v) \neq c'(v)\}$, the size of U' is even. Assume that $U' = \{u_1, u_2, \dots, u_{2k}\}$. For every $i \in [k]$ we consider a path from u_{2i-1} to u_{2i} and change signs to the opposite of variables corresponding to the edges of this path in the second Tseitin formula. Such operation for u_{2i-1} and u_{2i} corresponds to changing of the value of the labeling function in vertices $v \in u_{2i-1}$ and u_{2i} , labels for other vertices does not change since we change the signs for two edges. Applying this transformation for all connected components we transform the second Tseitin formula to the first. \square

Treewidth, pathwidth and minors. *A tree decomposition of an undirected graph $G(V, E)$ is a tree $T = (V_T, E_T)$ such that every vertex $u \in V_T$ corresponds to a set $X_u \subseteq V$ and it satisfies the following properties: 1. The union of X_u for $u \in V_T$ equals V . 2. For every edge $(a, b) \in E$ there exists $u \in V_T$ such*

that $a, b \in X_u$. 3. If a vertex $a \in V$ is in the sets X_u and X_v for some $u, v \in V_T$, then it is also in X_w for all w on the path between u and v in T .

If a tree T is a path, then this representation is a *path decomposition*. The *width of a tree decomposition* is the maximum $|X_u|$ for $u \in V_T$ minus one. A *treewidth* of a graph G is the minimal value of the width among all tree decompositions of the graph G . We denote it as $\text{tw}(G)$. The *pathwidth* of a graph G is the minimal value of the width among all path decompositions of a graph G . We denote it as $\text{pw}(G)$.

Lemma 2.6 ([13]). *For every graph G on n vertices $\text{pw}(G) = O(\log(n) \cdot \text{tw}(G))$.*

A *minor* of an undirected graph G is a graph that can be obtained from a graph G by a sequence of edge contractions, edge deletions and vertex deletions.

Theorem 2.7 ([5]). *For every constant $\delta < 1/36$ every graph G contains a $t \times t$ grid as a minor, where $t = \Omega(\text{tw}(G)^\delta)$.*

3 Lower bound for grids

In this section we prove the following Theorem.

Theorem 3.1. *Let T_n be an $n \times n$ grid graph. Then if a Tseitin formula $\text{TS}_{T_n, c}$ is satisfiable, then every 1-NBP that computes $\text{TS}_{T_n, c}$ has size at least $2^{\Omega(n)}$.*

Proof. T_n contains $(n+1)^2$ vertices and $2n(n+1)$ edges. In order to prove this theorem we use Lemma 2.4 for $l = n(n+1)$ (so l is the half of the number of edges). So we have to prove that if we delete half of the edges of T_n , then the resulting graph will have at most $\frac{(n+1)^2}{2} - \varepsilon \cdot n$ connected components for some constant $\varepsilon > 0$. Hence, by Lemma 2.4, every 1-NBP for $\text{TS}_{T_n, c}$ has size at least $2^{2\varepsilon n+1}$.

We call a subgraph of T_n *optimal* if it contains l edges and has the maximal number of connected components. The plan of the proof is the following. At first we show that there exists an optimal subgraph H that has one connected component that contains all edges and all other connected components are isolated vertices. Then we estimate the number of connected components of H .

Lemma 3.2. *There is an optimal subgraph of T_n that has exactly one connected component with at least two vertices.*

Proof. Consider all optimal subgraphs of T_n . Choose among them a subgraph H that contains a connected component M with the maximal number of edges. If M contains all edges of H , then the lemma is proved. Further we assume that not all edges are in M .

Consider the properties of the chosen graph H .

1. All the edges of the grid T_n between vertices of M are in M . Indeed, otherwise we can delete an edge from another connected component and add it to M . After this operation the number of connected components does not decrease, but the number of edges in M is strictly increased. This is a contradiction since M has the maximal number of edges among all the optimal subgraphs.

2. Every connected component, except M , is edge-biconnected (i.e. it is impossible to increase the number of connected components by the deletion of an edge from it). Indeed, assume that for some connected component except M it is possible to delete an edge from it such that the number of connected component increases. Then we delete this edge and add an edge of the grid that connects M with a vertex out of M . In this case the number of connected components is not changed but the number of edges in the maximal component would be increased. This is a contradiction.

3. There is no vertex v of T_n such that it is not in M but there are at least two edges between v and vertices of M in T_n . Proof by contradiction, assume that such a vertex exists. Consider a connected component K that differs from M and has edges. Consider a set of the lowest vertices in K and let u be the leftmost vertex among them. There are no edges to the left or down from the vertex u in the graph H . Since the connected component K has edges, there is at least one edge that is incident to u . By the

previous property, K is edge-biconnected, hence, u has precisely two incident edges. Let us delete the two edges incident to u from K and add two edges that connect the vertex v and M . The number of connected components doesn't decrease, but the number of edges in the maximal component increases. This is a contradiction.

4. Every 1×1 square of the grid T_n contains 0, 1 or 4 edges from M . A 1×1 square cannot contain exactly 3 edges because it contradicts the property 1. Let an 1×1 square contains exactly 2 edges from M . If these are two incident edges then we get a contradiction with the property 3 or the property 1. If these are two opposite edges, then we get a contradiction with the property 1.

5. For every $u, v \in M$, the minimal rectangle of the grid that contains both u and v (with all interior edges) is a subgraph of M (one of the sides of the rectangle could be of zero length, in that case it's just a line of the grid). It can be easily shown by the induction on the length of the shortest path between u and v using the property 4.

6. M is a rectangle of the grid with all edges of this rectangle. Consider the maximal rectangle of the grid that is fully contained (with all edges) in M . If there are vertices in M that are not in this rectangle, then we could increase this rectangle using the property 5.

So M is a rectangle of the grid. We say that M can be moved one step to the left (right, down or up) if all left (right, down or up) neighbours of the left (right, down or up) border of M are isolated vertices in H . Such a move doesn't change the number of connected components and the number of edges in them. Consider some connected component K that differs from M and contains edges. We move M one step closer to K in the way of decreasing the distance between them while it is possible. By the distance we understand the minimal L_1 -distance between two vertices from M and K . After some step it is not possible anymore; it means that one of the borders of M (w.l.o.g. it is the upper border) has upper neighbours from connected components that consist of more than one vertex.

Let M be a rectangle $x \times y$, where x, y are non-negative integers. That means that every horizontal line of M contains $x + 1$ vertices. Assume that among upper neighbours of the upper border of M there are m vertices that are in some connected component that consists of more than one vertex. Let these m vertices be in k connected components. Assume that there are r edges of the graph H between $x + 1$ upper neighbours of M (see an example on the left part of Figure 1). Since every edge between upper neighbours of M decreases the number of connected components, the following inequality holds $k \leq m - r$. Obviously, $r \leq x$.






Figure 1: Example: $x = 6, y = 4, r = 3, k = 3, m = 7$

Figure 2: Spines

Consider the following modification of the graph H : we move the rectangle M one step up and add edges down from r vertices on the bottom border of M (see example on the right part of Figure 1). The number of edges after this transformation is not changed since r edges overlapped and we added r edges. Now we estimate the number of connected components. On the bottom border we add $(x + 1) - r$ new connected components. On the upper border $(x + 1) - m + k$ connected components disappeared (were merged to one). Finally, the number of connected components increased by $m - k - r$, that is at least zero since $k \leq m - r$. But the number of edges in the maximal connected component increases, this contradicts the choice of the graph H . So we get a contradiction with the assumption that there are more than one connected components with at least one edge. \square

So we may assume that there is an optimal graph H that has one connected component M with at least two vertices and t isolated vertices. Notice that M is not necessary a rectangle now (in Lemma 3.2 we only

show that it has to be a rectangle under the assumption that the statement of the lemma is wrong). We are going to estimate $k_{T_n}(l) = t + 1$ from the above.

For every connected component K we define a set of *spines* that go out of it. We assume that T_n is a part of the infinite grid. An edge e of the infinite grid is a spine of K if it connects a vertex u from K with a vertex out of K and u is the right or bottom endpoint of e (see Figure 2). Note that spines can go outside of the square T_n if u is on the upper or left border of T_n . If a component is an isolated vertex, then it has exactly two spines.

The same edge cannot be a spine for two different connected components since we choose only the edges that go up or to the left from the component. Let h be the total number of all spines for all of the connected components of H . Every spine is either an edge of the square T_n that is not in H or is among $2(n + 1)$ edges that go outside the square T_n . Hence, $2(n + 1) + n(n + 1) \geq h$.

Let X be the number of spines of the connected component M , then we get that $h = 2t + X$. Using the previous inequality we estimate t as follows: $t \leq (n + 1)(n + 2)/2 - X/2 = (n + 1)^2/2 + (n + 1 - X)/2$. Hence, we need to show that $X \geq (1 + \epsilon)n$ for some constant ϵ .

Consider the minimal grid rectangle that contains M . Assume it has size $(a - 1) \times (b - 1)$, where a, b are natural numbers. Then there are spines of M in every of a vertical lines and in every of b horizontal lines. Then we get that $X \geq a + b \geq 2\sqrt{ab}$. On the other hand, the component M contains exactly $n(n + 1)$ edges, they need to be embedded into a rectangle $(a - 1) \times (b - 1)$ that contains $a(b - 1) + b(a - 1)$ edges. Then we can estimate $2ab > a(b - 1) + b(a - 1) \geq n(n + 1)$ and we get $X \geq \sqrt{2n(n + 1)} > \sqrt{2}n$.

Using the upper bound on t , the number of connected components can be estimated as: $k_{T_n}(l) = t + 1 \leq (n + 1)^2/2 - (\sqrt{2} - 1)n + \frac{3}{2}$. Then by Lemma 2.4, every 1-NBP for a satisfiable Tseitin formula $\text{TS}_{T_n, c}$ has size at least $2^{2(\sqrt{2}-1)n-2}$. \square

4 Treewidth

The main goal of this section is to prove the following theorem.

Theorem 4.1. *Let $\text{TS}_{G, c}$ be a satisfiable Tseitin formula. Then every 1-NBP for $\text{TS}_{G, c}$ has size at least $2^{\Omega(\text{tw}(G)^\delta)}$ for all $\delta < 1/36$.*

Lemma 4.2. *Let D be a 1-NBP computing a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. If we change every node in D labeled with the variable x_1 by a guessing node and remove all labels of all its outgoing edges, then we obtain a valid 1-NBP that computes $\exists x_1 f(x_1, x_2, \dots, x_n)$.*

Proof. For every path from the source to a sink in D' every variable occurs at most once, so D' is a valid 1-NBP. Consider an assignment τ to the variables x_2, \dots, x_n . Let us consider a path in D' from the source to the sink labeled with 1 that is consistent with τ . Consider the same path in D , it also finishes in the sink labeled with 1. Then there exists an assignment τ' that satisfies f and τ' extends τ on a value of the variable x_1 . Now consider a satisfying assignment σ of the function $\exists x_1 f(x_1, x_2, \dots, x_n)$. For some value of x_1 there exists a path from the source to the sink labeled with 1 in the branching program D that is consistent with this value and with σ . The copy of this path in the branching D' is consistent with σ and reaches the sink labeled with 1. \square

Lemma 4.3. *Let H be a minor of an undirected graph G and Tseitin formulas $\text{TS}_{G, c}$ and $\text{TS}_{H, c'}$ be satisfiable. Then for every S and every 1-NBP of size S that computes $\text{TS}_{G, c}$, there is a 1-NBP that computes $\text{TS}_{H, c'}$ of size at most S .*

Proof. It suffices to prove the statement of the lemma for the case when H is obtained from G by the application of one operation. Let us consider all types of operations separately.

1. H is obtained from G by the deletion of an edge e . Let σ be a satisfying assignment of $\text{TS}_{G, c}$. We apply the partial assignment $x_e := \sigma(x_e)$ to $\text{TS}_{G, c}$ and get the satisfiable formula $\text{TS}_{H, c''}$. It is well known that the application of a substitution does not increase size of the 1-NBP. By Lemma 2.5, sizes of the minimal 1-NBPs for $\text{TS}_{H, c''}$ and $\text{TS}_{H, c'}$ are equal.

2. H is obtained from G by the deletion of a vertex v . Since all the variables of Tseitin formulas are associated with edges, this case can be considered as a sequence of edge deletions.

3. A graph $H(V_H, E_H)$ is obtained from a graph $G(V, E)$ by the contraction of an edge $e = (u, v)$. Let us define a labeling function $c'' : V_H \rightarrow \{0, 1\}$ as follows: for all the vertices that differ from the joined vertex $\{u, v\}$ it has the same value as in the labeling function c and $c''(\{u, v\}) = c(u) + c(v) \bmod 2$. By the construction, every connected component U of H corresponds to a connected component U' of G with the same sum of the labeling functions: $\sum_{w \in U} c''(w) = \sum_{w \in U'} c(w) \bmod 2$. Hence, by Lemma 2.2, $\text{TS}_{H, c''}$ is satisfiable.

Lemma 4.4. *Formulas $\text{TS}_{H, c''}$ and $\exists x_e \text{TS}_{G, c}$ define the same function.*

Proof. Formulas $\text{TS}_{H, c''}$ and $\exists x_e \text{TS}_{G, c}$ depended on the same set of variables $X = \{x_l \mid l \in E \setminus \{e\}\}$, since H is obtained from G by the contraction of the edge e . Consider some assignment of these variables $\sigma : X \rightarrow \{0, 1\}$. Assume that σ satisfies $\text{TS}_{H, c''}$. Consider an assignment $\sigma' : X \cup \{x_e\} \rightarrow \{0, 1\}$ such that for every $x \in X$, $\sigma'(x) = \sigma(x)$ and $\sigma'(x_e) = \sum_{t \in E_u} \sigma(x_t) + c(u) \bmod 2$, where E_u is a set of edges that are incident to the vertex u in G except for the edge e . We need to check that σ' satisfies $\text{TS}_{G, c}$. All conditions that are not related to vertices u and v are satisfied automatically, because they do not contain the edge e . The condition in the vertex u is satisfied because of the choice of the value of the variable x_e in σ' . Let E_v be a set of edges that are incident to the vertex v in G except for the edge e . To check the condition in the vertex v we should compute $\sigma'(x_e) + \sum_{t \in E_v} \sigma(x_t)$. By the definition of the value $\sigma'(x_e)$ we get that the sum equals $\sum_{t \in E_u \cup E_v} \sigma(x_t) + c(u)$ that is equal to $c(v)$, because σ satisfies the condition in the vertex $\{u, v\}$ in the formula $\text{TS}_{H, c''}$.

Now let σ be an assignment that satisfies $\exists x_e \text{TS}_{G, c}$. σ can be extended to an assignment σ' that satisfies $\text{TS}_{G, c}$. Let us check that σ satisfies $\text{TS}_{H, c''}$. We should check the condition only for the vertex $\{u, v\}$: $\sum_{t \in E_u \cup E_v} \sigma(x_t) = (\sum_{t \in E_u} \sigma(x_t) + \sigma'(x_e)) + (\sum_{t \in E_v} \sigma(x_t) + \sigma'(x_e)) = c(u) + c(v)$. \square

By Lemma 4.2, the minimal size of a 1-NBP for $\exists x_e \text{TS}_{G, c}$ (that is by Lemma 4.4 equivalent to $\text{TS}_{H, c''}$) is at most the minimal size of a 1-NBP for $\text{TS}_{G, c}$. By Lemma 2.5, minimal sizes of 1-NBPs for $\text{TS}_{H, c''}$ and for $\text{TS}_{H, c'}$ are equal. \square

Proof of Theorem 4.1. By Theorem 2.7, the graph G contains a $t \times t$ grid graph as a minor, where $t = \Omega(\text{tw}(G)^\delta)$. The theorem follows from Theorem 3.1 and Lemma 4.3. \square

We also obtain an upper bound:

Theorem 4.5. *Every satisfiable Tseitin formula $\text{TS}_{G, c}$ can be represented as OBDD of size $|E|2^{\text{pw}(G)+1} + 2$.*

Proof. Consider a path decomposition of the graph G : X_1, X_2, \dots, X_s such that $|X_i| \leq \text{pw}(G) + 1$ for $i \in [s]$. We split the set of edges into disjoint (maybe empty) parts: E_1, E_2, \dots, E_s . E_1 is a set of edges between vertices X_1 , if $i > 1$, then E_i is a set of edges between vertices in X_i that did not appear earlier (it is enough to require that no one of these edges is between vertices from X_{i-1} , since a vertex can not disappear and then appear again). Since every edge occurs in at least one X_i , every edge will be in some set E_j .

We fix some order of edges in which the first are edges from E_1 , then edges from E_2 , \dots and in the end there are edges from E_s . Let e_i be i th edge in this order.

Consider a complete decision tree T that computes $\text{TS}_{G, c}$ and such that in every path from the root of T to a leaf variables are requested in the order: $e_1, e_2, \dots, e_{|E|}$. We define the following equivalence relation on the nodes of T : two nodes of T are equivalent if partial substitutions corresponding to them either falsify a clause of $\text{TS}_{G, c}$ or equally modify $\text{TS}_{G, c}$. Nodes of the constructing OBDD are the equivalence classes of the defined relation. The source is the equivalent class of the root of T , the sink labeled with 0 is the equivalent class of nodes of T that falsify a clause of $\text{TS}_{G, c}$ and the sink labeled with 1 is the equivalent class of leaves of T labeled with 1. It is easy to verify that two nodes of T from the same equivalence class except the sink 0 have the same distance from the root and they are labeled with the same edge of G . A label of an equivalence class in OBDD is the label of any of its representative. If (u, v) is an edge of T labeled with $a \in \{0, 1\}$, then we introduce an edge labeled with a between equivalence classes of u and v . It is straightforward that we get

a correct OBDD that computes $\text{TS}_{G,c}$. Consider a level l of T and let us estimate the number of equivalent classes of vertices from the level l (except the sink 0). Let $e_l \in E_j$. Consider a vertex $v \in V \setminus X_j$ and let $v \in X_i$. If $i < j$, then all incident edges of v are among e_1, e_2, \dots, e_l , hence the parity condition in the vertex v is either falsified by the substitution (and then the current node is from the equivalence class of the sink 0) or is satisfied by the substitution. If $i > j$, then edges e_1, e_2, \dots, e_l are not incident to v . Thus equivalence classes for nodes of T (except the sink 0) on the level l differ only on how they modify the labeling function for vertices from X_i . Since $|X_i| \leq \text{pw}(G) + 1$, then the number of equivalence classes for nodes on level l (except the sink 0) is at most $2^{\text{pw}(G)+1}$. There are $|E|$ levels and we also have two sinks, thus size of the constructed OBDD is at most $|E|2^{\text{pw}(G)+1} + 2$. \square

Corollary 4.6. *Any satisfiable Tseitin formula based on a graph $G(V, E)$ can be represented as OBDD of size $O(|E||V|^{O(\text{tw}(G))})$.*

Proof. Follows from Theorem 4.5 and Lemma 2.6. \square

5 Lower bound in the proof system $\text{OBDD}(\wedge, \text{reordering})$

In this section we show that any refutation of an unsatisfiable Tseitin formula $\text{TS}_{G,c}$ in the proof system $\text{OBDD}(\wedge, \text{reordering})$ has size at least $2^{\Omega(\text{tw}(G)^\delta)}$ for all $\delta < 1/36$.

If F is a formula in CNF, we say that a sequence D_1, D_2, \dots, D_t of OBDDs is an $\text{OBDD}(\wedge, \text{reordering})$ -refutation of F if D_t is an OBDD that represents the constant false function, and for all $1 \leq i \leq t$, D_i is an OBDD that represents a clause of F or can be obtained from the previous D_j 's by one of the following inference rules: (conjunction or join) D_i is an OBDD with order π , that represents $D_k \wedge D_l$ for $1 \leq l, k < i$, where D_k, D_l have the same order π ; (reordering) D_i is an OBDD that is equivalent to an OBDD D_j with $j < i$ (note that D_i and D_j may have different orders).

We say that a graph H is *t-good* if H is connected and every OBDD-representation of any satisfiable Tseitin formula $\text{TS}_{H,c}$ has size at least t . The following theorem can be proved using ideas from [12]:

Theorem 5.1. [cf. [12]] *Let $G(V, E)$ be a connected graph and degrees of all vertices of G be bounded by a constant. Assume that the graph G has the following properties: 1) if we delete any vertex from G , we get a *t-good* graph; 2) for every two vertices u and v of G there is a path p between them such that if we delete all vertices from p , we get a *t-good* graph. And if we delete from G vertices u and v and the edges of the path p , we also get a connected graph.*

Then any $\text{OBDD}(\wedge, \text{reordering})$ -refutation of an unsatisfiable Tseitin formula $\text{TS}_{G,c}$ has size at least $\Omega(\sqrt{t})$.

We start with the idea of the proof. We consider the last step of the $\text{OBDD}(\wedge, \text{reordering})$ -refutation: the conjunction of OBDDs F_1 and F_2 is the identically false function but both F_1 and F_2 are satisfiable. Both F_1 and F_2 are conjunctions of several clauses of $\text{TS}_{G,c}$.

Since G remains connected after removing of any single vertex, F_1 and F_2 together contain all clauses of $\text{TS}_{G,c}$. Assume that there are two nonadjacent vertices u and v such that F_1 does not contain a clause C_u that corresponds to the vertex u and F_2 does not contain a clause C_v that corresponds to v (if this assumption is false, the proof is rather straightforward). We consider two partial substitutions ρ_1 and ρ_2 that are both defined on the edges adjacent to u and v and on the edges of the path p between u and v . The substitutions ρ_1 and ρ_2 assign opposite values to edges of the path p and are consistent on all other edges. The substitution ρ_1 satisfies C_v and refutes C_u and ρ_2 satisfies C_u and refutes C_v .

By the construction $F_1|_{\rho_1} \wedge F_2|_{\rho_2}$ is almost a satisfiable Tseitin formula based on the graph that is obtained from G by deletion of the vertices u and v and all edges from the path p . However, it is also possible that this formula does not contain some clauses for the vertices from p . Thus we make additional partial substitution τ that substitute values from a satisfying assignment for all remaining edges for vertices from p . $(F_1|_{\rho_1} \wedge F_2|_{\rho_2})|_{\tau}$ is satisfiable Tseitin formula based on the graph that is obtained from G by deletion of all vertices from the path p . The size of an OBDD representation of such a formula is at least t by the condition of the theorem. Hence by Lemma 2.1 we get that either F_1 or F_2 has size at least $\Omega(\sqrt{t})$ in the given order.

Proof of Theorem 5.1. Consider the last step of the proof: conjunction of OBDDs F_1 and F_2 is the identically false function but F_1 and F_2 are satisfiable. Both F_1 and F_2 are conjunctions of several clauses of $\text{TS}_{G,c}$. Every clause of $\text{TS}_{G,c}$ is either in F_1 or in F_2 since otherwise $F_1 \wedge F_2$ is satisfiable by Lemma 5.2. Our goal is to prove that either F_1 or F_2 has size at least \sqrt{t} .

Lemma 5.2. *Suppose Φ denotes the formula that can be obtained from $\text{TS}_{G,c}$ by removing one arbitrary clause. Then Φ is satisfiable and every OBDD representation of Φ has size at least t .*

Proof. Assume that $\text{TS}_{G,c} = \Phi \wedge C$, where C is a clause that corresponds to the equation in a vertex v . Let H be the result of the removing of the vertex v from G . H is t -good and hence graph H is connected. Suppose we make a substitution to all variables of the clause C that falsifies C . Let Φ' denote the result of this substitution applied to Φ . Since the substitution falsifies the parity condition at the vertex v , Φ' corresponds to a satisfiable Tseitin formula based on H , hence the size of every OBDD for Φ' is at least t . Φ' is the result of the substitution applied to Φ , hence by Lemma 2.1 the size of every OBDD representing Φ is at least t . \square

We consider two cases:

1. There exist two non-adjacent vertices u and v from G such that F_1 does not include some clause C_v that corresponds to vertex v and F_2 does not include some clause C_u for vertex u .

Consider a path p from v to u that exists by the property 2 from the conditions of the theorem. Let e_v be the first edge of p and e_u be the last edge ($e_v \neq e_u$ since u and v are non-adjacent). Consider two substitutions ρ_1 and ρ_2 with the same support: all edges that are incident to vertices u and v and all edges from the path p . Substitutions ρ_1 and ρ_2 are consistent on edges that are out of p : all edges that are adjacent to u or v but not in p have values that do not satisfy C_u and C_v (this is possible since u and v are non-adjacent). ρ_1 substitutes zeros to all edges from p except e_u and e_v and substitute a value to e_v that does not satisfy C_v and a value to e_u that satisfies C_u . ρ_2 substitute ones to all edges from p except e_u and e_v and substitute a value to e_v that satisfies C_v and a value to e_u that does not satisfy C_u . So edges from p have different values in ρ_1 and ρ_2 ; ρ_1 satisfies u and refutes v and ρ_2 refutes u and satisfies v .

Consider the graph G' that can be obtained from G by removing u , v and all edges from the path p . The graph G' is connected by the property 2 from the conditions of the theorem.

Let c' be a labeling function of the result of the substitution ρ_1 applied to $\text{TS}_{G,c}$ and c'' be a restriction of c' on $V \setminus \{u, v\}$. Note that ρ_2 corresponds to the same c'' since ρ_1 and ρ_2 identically change the value of the labelling function for all vertices except u and v . We claim that $\text{TS}_{G',c''}$ is satisfiable. Indeed if we make a substitution ρ_1 to $\text{TS}_{G,c}$ the vertex v would be refuted, the vertex u would be satisfied, all other vertices are marked according c' . Thus the sum of values of c'' is even and $\text{TS}_{G',c''}$ is satisfiable since G' is connected.

We consider the conjunction $F_1|_{\rho_1} \wedge F_2|_{\rho_2}$. Any satisfying assignment of $\text{TS}_{G',c''}$ satisfies both $F_1|_{\rho_1}$ and $F_2|_{\rho_2}$, hence $F_1|_{\rho_1} \wedge F_2|_{\rho_2}$ is satisfiable. Suppose we represent $F_1|_{\rho_1} \wedge F_2|_{\rho_2}$ as a conjunction of clauses. For every vertex w that is not in p , the union of clauses of F_1 and F_2 contains all clauses that correspond to the equation at vertex w , substitutions ρ_1 and ρ_2 are consistent for all variables from this equation, hence $F_1|_{\rho_1} \wedge F_2|_{\rho_2}$ contains all clauses that correspond to the equation of vertex w in formula $\text{TS}_{G',c''}$. Consider a partial substitution τ that substitutes values to edges of G' that are incident to vertices of p according some satisfying assignment of $\text{TS}_{G',c''}$. If we delete all vertices of the path p from G' we get a graph G'' that is connected by the property 2 of conditions of the theorem. Since τ is consistent with a satisfying assignment of $\text{TS}_{G',c''}$, then the application of the substitution τ converts $\text{TS}_{G',c''}$ to a satisfiable Tseitin formula $\text{TS}_{G'',c''}$. And $\text{TS}_{G'',c''}$ coincides with $(F_1|_{\rho_1} \wedge F_2|_{\rho_2})|_{\tau}$.

Size of any OBDD representation of the formula $\text{TS}_{G'',c''}$ is at least t since G'' is t -good. Then any OBDD representing $(F_1|_{\rho_1} \wedge F_2|_{\rho_2})|_{\tau}$ has size at least t , hence by Lemma 2.1 any OBDD representing

$F_1|_{\rho_1} \wedge F_2|_{\rho_2}$) has size at least t . Thus, by Lemma 2.1 for every given order of variables π either $F_1|_{\rho_1}$ or $F_2|_{\rho_2}$ has size of the minimal π -OBDD at least \sqrt{t} . Hence, the minimal π -OBDD for F_1 or F_2 has size at least \sqrt{t} .

2. In the second case there are no such non-adjacent vertices. Since F_1 is not identically false, there exists a vertex u such that F_1 does not include a clause that corresponds to a vertex u and by the assumption F_2 does not include clauses only for the vertex v and its neighbours. Since G is a constant-degree graph, F_2 differs from a Tseitin formula without one clause by the constant number of clauses that depend on the constant number of variables. Any function that depends only on the constant number of variables may be represented by OBDD of constant size. Thus any OBDD representation of F_2 by Lemmas 5.2 and 2.1 has size at least $\Omega(t)$.

□

Theorem 5.1 is used in the proof of the main result of this section:

Theorem 5.3. *Let G be a connected graph and $\text{TS}_{G,c}$ be an unsatisfiable formula. Then every OBDD(\wedge , reordering)-refutation of $\text{TS}_{G,c}$ has size at least $2^{\Omega(\text{tw}(G)^\delta)}$ for all $\delta < 1/36$.*

We will need the following lemmas.

Lemma 5.4. *Let ϕ be an unsatisfiable CNF formula that has a refutation in the proof system OBDD(\wedge , reordering) of size S . Let ρ be a partial substitution of values of the formula ϕ . Then $\phi|_\rho$ has an OBDD(\wedge , reordering)-refutation of size at most S .*

Proof. Let D_1, D_2, \dots, D_s be a refutation of the formula ϕ , then $D_1|_\rho, D_2|_\rho, \dots, D_s|_\rho$ is a refutation of the formula $\phi|_\rho$, where $D_i|_\rho$ is a result of a substitution ρ to OBDD D_i . By Lemma 2.1, the application of a substitution does not increase the size of OBDD. □

Lemma 5.5. *Let $G(V, E)$ be a connected graph and $G'(V', E')$ be a connected subgraph of G with $E' \neq \emptyset$ that is obtained from G by the deletion of some vertices and edges. For every unsatisfiable Tseitin formula $\text{TS}_{G,c}$ there exists a substitution ρ on variables $E \setminus E'$, such that ρ does not falsify any clause of $\text{TS}_{G,c}$.*

Proof. We substitute one-by-one values of variables that correspond to edges from $E \setminus E'$ (and correspondingly modify the Tseitin formula) and we maintain the following invariant while deleting edges of G : a Tseitin formula corresponding to every connected component that doesn't contain vertices from V' is satisfiable. The invariant holds at the beginning, since the graph G' is connected and thus there are no such connected components. Suppose we assign an edge e , if e is from an unsatisfiable connected component (i.e. a part of Tseitin formula corresponding to this component is unsatisfiable) and it is not a bridge of G , then we assign it with the value 0, after this transformation the component stays unsatisfiable. If e is a bridge of an unsatisfiable component, then we assign it with a value such that the new connected component that is obtained after the deletion of e and that does not contain vertices from V' would be satisfiable (such value exists by Lemma 2.2). If e is from a satisfiable connected component, then we assign a value from the satisfying assignment of this connected component. In the end the invariant implies that the resulting substitution does not falsify any clause of $\text{TS}_{G,c}$. □

Proof of Theorem 5.3. By Theorem 2.7 the graph G contains a $s \times s$ grid minor H , where $s = \Omega(\text{tw}(G)^\delta)$. Consider a process of obtaining the minor H from the graph G such that this process contains the minimal possible number of operations of edge contractions. Since operations of edge and vertex deletions and edge contractions commute, we denote by G' a graph that is obtained from G after the application of all edge and vertex deletions. Then H can be obtained from G' by the application of only edge contractions. We assume that if during a step of edge contractions we get parallel edges, then all of them except one should be deleted.

In Lemmas 5.6, 5.7, 5.8 and 5.9 we verify that G' satisfies the conditions of Theorem 5.1. And then we conclude the proof by the application of Lemma 5.5 and Lemma 5.4.

Lemma 5.6. *All vertices in the graph G' have degree at most 4.*

Proof. Note that after the application of an edge contraction the degree of a vertex cannot decrease, since otherwise it is possible to replace an edge contraction by a vertex deletion. Since all vertices in H have degrees at most 4, it must also be true for G' . \square

Lemma 5.7. *After the deletion of any vertex from the graph G' we get a connected graph.*

Proof. Suppose that we delete a vertex v from the graph G' and get a graph with at least two connected components K_1 and K_2 . We know that we can obtain the $s \times s$ grid H from the graph G' by several edge contractions. Suppose the vertex v is transformed to the vertex v' of H after the application of these edge contractions. Since H remains connected after the deletion of the vertex v' , it is not possible that some vertex $a \in K_1$ and some vertex $b \in K_2$ are transformed to vertices a' and b' from $H - v'$. Indeed, $H - v'$ is connected, hence there is a path in H connecting a' and b' that does not contain v' , hence there is a path in G' connecting a and b that does not contain v ; this is a contradiction since a and b are from different connected components of $G' - v$. Thus for some $i \in \{1, 2\}$ the component K_i is fully transformed to the vertex v' . It means that we could just delete all vertices of K_i and do not contract anything, it decreases the number of contractions and contradicts the choice of the graph G' . \square

Consider the $s \times s$ grid square H and for every two vertices u and v of H we define a path $p'_{u,v}$ between them. If at least one of u and v is not on the border of the square H , then $p'_{u,v}$ consists of two parts connected at an endpoint, one of whose is horizontal and the other is vertical. If u and v belong to the border of H , then $p'_{u,v}$ is the shortest path between u and v that only passes through vertices on the border. In both cases it can be easily verified that the $s \times s$ graph remains connected after the deletion of vertices u and v and edges of $p'_{u,v}$, and also after the deletion of all vertices of $p'_{u,v}$. Notice the following property: if two vertices of $p'_{u,v}$ are connected by an edge of the grid, then they are consequent in the path $p'_{u,v}$.

Consider two arbitrary vertices u and v of the graph G' , suppose that after the application of edge contractions they are transformed to the vertices u' and v' (u' and v' might coincide) of the $s \times s$ grid H . Let $p_{u,v}$ be a path between u and v in the graph G' that is transformed to $p'_{u',v'}$. If $u = v$, we assume that $p_{u,u}$ is a path with the only one vertex u .

Lemma 5.8. *For every two vertices u, v of the graph G' if we delete u and v and edges of $p_{u,v}$ from G' we obtain a connected graph. And if we delete all of the vertices of $p_{u,v}$ we also obtain a connected graph.*

Proof. At first we note that it is impossible that G' has two adjacent vertices a and b out of the path $p_{u,v}$ such that a and b are transformed to different pair of distinct vertices a' and b' of path $p'_{u',v'}$. Indeed, since we do not delete edges, a' and b' are connected by an edge in H , hence a' and b' are two consecutive vertices of $p'_{u',v'}$. Since $p_{u,v}$ is transformed to $p'_{u',v'}$, there are vertices s and t from $p_{u,v}$ that are transformed to a' and b' . Vertices s and t are connected by edges of $p_{u,v}$, hence there is an edge of $p_{u,v}$ that is transformed to an edge between a' and b' . The edge between a and b is also transformed to an edge between a' and b' . This is a contradiction, since we get two edges between a' and b' that come from different edges of G' , hence one of them should be deleted but we assume that there are no more deletions.

We claim that the graph G' remains connected if we delete from it vertices u, v and edges of the path $p_{u,v}$. Proof by contradiction. Suppose that after the deletion we will get a graph with at least two connected components K_1 and K_2 . Since H remains connected after deletions of vertices u', v' and edges of the path $p'_{u',v'}$, then it is not possible that some vertices $a \in K_1$ and $b \in K_2$ are transformed to vertices a', b' from $H - \{v', u'\}$, because they are connected in $H - \{v', u'\}$ and thus there is a path between them that does not go through vertices u, v and edges of the path $p_{u,v}$ in the initial graph G' . Thus, for some $i \in \{1, 2\}$ every vertex of the component K_i is transformed to v' or to u' . We have already shown that it is impossible that some vertices from K_i are transformed to v' and some to u' (otherwise there are two adjacent vertices with such property). Hence we may just delete all vertices from K_i and do not contract anything and this decreases the number of contractions. And we get a contradiction with the choice of the graph G' .

By similar arguments the graph G' remains connected if we delete from it all vertices of $p_{u,v}$. \square

Lemma 5.9. *For every two vertices u, v of the graph G' the graph that is obtained from G' by the deletion of all vertices of $p_{u,v}$ contains a $\lfloor s/3 \rfloor \times \lfloor s/3 \rfloor$ grid as a minor.*

Proof. Consider some vertices u and v of the $s \times s$ square. We show that if we delete all vertices of the path $p'_{u,v}$ from the $s \times s$ square, then there is a sub-square of the size $\lfloor s/3 \rfloor \times \lfloor s/3 \rfloor$ that does not contain any deleted edge or vertex. Indeed, there are 9 squares of the size $\lfloor s/3 \rfloor \times \lfloor s/3 \rfloor$ that are fully contained in the square $s \times s$ and do not have common 1×1 squares. If the path goes only by the border of the square then it does not touch the central square $\lfloor s/3 \rfloor \times \lfloor s/3 \rfloor$. In the other case the path has two straight segments (may be of the zero length): vertical and horizontal. It is easy to see that such paths can have common vertices with at most 8 chosen squares, hence there is at least one $\lfloor s/3 \rfloor \times \lfloor s/3 \rfloor$ square that is not touched.

Suppose we deleted the path $p_{u,v}$ between two vertices u and v from the graph G' . We find in the square H a square $T_{\lfloor s/3 \rfloor}$ of the size $\lfloor s/3 \rfloor \times \lfloor s/3 \rfloor$ such that it does not intersect with vertices u', v' and vertices of the path $p'_{u',v'}$. Let U be the set of vertices of the graph G' that is contracted to vertices of $T_{\lfloor s/3 \rfloor}$. The set of vertices U does not intersect with u, v and vertices of the path $p_{u,v}$, since in the other case the square $T_{\lfloor s/3 \rfloor}$ would also intersect u', v' or p' . Thus $T_{\lfloor s/3 \rfloor}$ is a minor of a subgraph G' that is induced by vertices from U , and, hence, a minor of the graph G' . \square

Let K be a graph obtained from G' by the deletion of a vertex u or by the deletion of vertices from a path $p_{u,v}$ for some u, v . Lemma 5.9 and Lemma 4.3 imply that size of 1-NBP for every satisfiable Tseitin formula $TS_{K,c}$ is at least the minimal size of 1-NBP for a satisfiable Tseitin formula $TS_{T_{\lfloor s/3 \rfloor},c'}$ that by Theorem 3.1 is at least $2^{\Omega(s)}$. Using this and Lemmas 5.6,5.7,5.8 we get that G' satisfies the condition of Theorem 5.1 for $t = 2^{\Omega(s)}$.

By Lemma 5.5 there exists a substitution ρ of variables that correspond to edges of the graph G that are not in G' that does not falsify clauses of the formula $TS_{G,c}$. The application of ρ to the formula $TS_{G,c}$ transforms it to the formula $TS_{G',c'}$ (other clauses are satisfied). By Theorem 5.1 the size of every OBDD(\wedge ,reordering)-refutation of the formula $TS_{G,c}$ is $2^{\Omega(s)}$. Then by Lemma 5.4 the size of every OBDD(\wedge ,reordering)-refutation of the formula $TS_{G,c}$ is $2^{\Omega(s)}$. \square

Acknowledgements. The authors are grateful to Fedor Fomin and Alexander Kulikov for fruitful discussions and to Alexander Knop, Dmitry Sokolov and anonymous reviewers for useful comments. The second author is a Young Russian Mathematics award winner and would like to thank sponsors and jury of the contest.

The research was supported by Russian Science Foundation (project 16-11-10123).

References

- [1] Michael Alekhnovich and Alexander A. Razborov. Satisfiability, branch-width and tseitin tautologies. *Computational Complexity*, 20(4):649–678, 2011.
- [2] Albert Atserias, Phokion G. Kolaitis, and Moshe Y Vardi. Constraint propagation as a proof system. In *CP 2004*, volume 3258 of *LNCS*, pages 77–91. Springer, 2004.
- [3] Sam Buss, Dmitry Itsykson, Alexander Knop, and Dmitry Sokolov. Reordering rule makes OBDD proof systems stronger. In *CCC-2018*, pages 16:1–16:24, 2018.
- [4] Chandra Chekuri and Julia Chuzhoy. Polynomial bounds for the grid-minor theorem. *J. ACM*, 63(5):40:1–40:65, 2016.
- [5] Julia Chuzhoy. Excluded grid theorem: Improved and simplified. In *Proceedings of STOC-2015*, pages 645–654, 2015.
- [6] Stefan S. Dantchev and Søren Riis. "planar" tautologies hard for resolution. In *FOCS-2001*, pages 220–229, 2001.

- [7] Andrea Ferrara, Guoqiang Pan, and Moshe Y. Vardi. Treewidth in verification: Local vs. global. In *LPAR-2005*, pages 489–503. Springer, 2005.
- [8] Nicola Galesi, Navid Talebanfard, and Jacobo Torán. Cops-robber games and the resolution of tseitin formulas. In *SAT-2018*, pages 311–326, 2018.
- [9] L. Glinskii and D. Itsykson. Satisfiable Tseitin formulas are hard for nondeterministic read-once branching programs. In *MFCS-2017*, pages 26:1–26:12, 2017.
- [10] Qian-Ping Gu and Hisao Tamaki. Improved bounds on the planar branchwidth with respect to the largest grid minor size. *Algorithmica*, 64(3):416–453, Nov 2012.
- [11] Johan Håstad. On small-depth Frege proofs for Tseitin for grids. In *FOCS-2017*, pages 97–108, 2017.
- [12] Dmitry Itsykson, Alexander Knop, Andrei Romashchenko, and Dmitry Sokolov. On OBDD-based algorithms and proof systems that dynamically change order of variables. In *STACS-2017*, pages 43:1–43:14, 2017.
- [13] Ephraim Korach and Nir Solel. Tree-width, path-width, and cutwidth. *Discrete Applied Mathematics*, 43(1):97–101, 1993.
- [14] Jan Krajíček. *Bounded arithmetic, propositional logic and complexity theory*. Cambridge University Press, 1995.
- [15] László Lovász, Moni Naor, Ilan Newman, and Avi Wigderson. Search Problems in the Decision Tree Model. *SIAM J. Discrete Math.*, 8(1):119–132, 1995.
- [16] Igor Razgon. On the read-once property of branching programs and cnfs of bounded treewidth. *Algorithmica*, 75(2):277–294, June 2016.
- [17] N. Robertson, P. Seymour, and R. Thomas. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 62(2):323 – 348, 1994.
- [18] Neil Robertson and Paul D. Seymour. Graph minors. v. excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986.
- [19] G. S. Tseitin. On the complexity of derivation in the propositional calculus. *Zapiski nauchnykh seminarov LOMI*, 8:234–259, 1968.
- [20] A. Urquhart. Hard examples for resolution. *JACM*, 34(1):209–219, 1987.
- [21] Ingo Wegener. *Branching Programs and Binary Decision Diagrams*. SIAM, 2000.