THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Exponential Lower Bounds for AC⁰-Frege Imply Superpolynomial Frege Lower Bounds

Citation for published version:

Filmus, Y, Pitassi, T & Santhanam, R 2011, Exponential Lower Bounds for AC⁰-Frege Imply Superpolynomial Frege Lower Bounds. in L Aceto, M Henzinger & J Sgall (eds), *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Proceedings, Part I*. Lecture Notes in Computer Science, vol. 6755, Springer, Zurich, Switzerland, pp. 618-629. https://doi.org/10.1007/978-3-642-22006-7_52

Digital Object Identifier (DOI):

[10.1007/978-3-642-22006-7_52](https://doi.org/10.1007/978-3-642-22006-7_52)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Proceedings, Part I

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Exponential Lower Bounds for AC^0 -Frege Imply Superpolynomial Frege Lower Bounds

Yuval Filmus¹ *, Toniann Pitassi¹ *, and Rahul Santhanam²

¹ Univeristy of Toronto, `yuvalf,toni@cs.toronto.edu`

² University of Edinburgh, `rsanthan@inf.ed.ac.uk`

Abstract. We give a general transformation which turns polynomial-size Frege proofs to subexponential-size AC^0 -Frege proofs. This indicates that proving exponential lower bounds for AC^0 -Frege is hard, since it is a longstanding open problem to prove super-polynomial lower bounds for Frege. Our construction is optimal for tree-like proofs.

As a consequence of our main result, we are able to shed some light on the question of weak automatizability for bounded-depth Frege systems. First, we present a simpler proof of the results of Bonnet et al. [5] showing that under cryptographic assumptions, bounded-depth Frege proofs are not weakly automatizable. Secondly, we show that because our proof is more general, under the right cryptographic assumptions, it could resolve the weak automatizability question for lower depth Frege systems.

* Supported by NSERC

1 Introduction

The fundamental question in computational complexity is the P vs. NP question. Though we are very far from resolving this question, over the past few decades we have made substantial progress in understanding why certain approaches, for example diagonalization and the use of combinatorial or algebraic techniques to prove circuit lower bounds, are unlikely to work. Various barriers such as the relativization, natural proofs and algebrization barriers have been formulated to capture the limitations of known techniques, and in turn, this meta-level understanding of complexity lower bound problems has led to developments in areas such as low-level complexity and derandomization.

However, there are still approaches whose power is not well understood, such as those used in proof complexity. Proof complexity was introduced by Cook and Reckhow [8] as a framework within which to study the NP vs. coNP problem. Cook and Reckhow defined *propositional proof systems* in a very general way by insisting only that proofs be verifiable in polynomial time, and showed that the existence of a propositional proof system in which all tautologies have polynomial size proofs is equivalent to $NP = coNP$. They suggested a program to separate NP and coNP (and thereby P and NP) by showing superpolynomial proof size lower bounds for explicit tautologies in progressively stronger proof systems. The hope was that techniques from logic and proof theory could be effective where techniques inspired by recursion theory or combinatorics are not. The fact that the very definition of the P vs. NP question involves the notion of “proof” in a fundamental way makes this hope somewhat plausible.

Indeed, over the past couple of decades, lower bounds have been shown for various natural proof systems [9, 3]. However, lower bounds for natural systems such as Frege and Extended Frege still seem out of reach. We seem to have hit a “wall” with proof complexity lower bounds, just as with circuit complexity lower bounds.

To an extent, this reflects the fact that the techniques used for the currently strongest proof complexity lower bounds are adaptations of the techniques used in circuit complexity, and limitations of the circuit complexity techniques carry over to the versions used in proof complexity. There is an informal “mapping” from proof systems to complexity classes, where a proof system Q corresponds to the smallest complexity class C such that the lines of polynomial-sized proofs in Q are functions in C. In this way, Resolution maps to DNFs, Bounded-Depth Frege to non-uniform AC^0 , Frege to non-uniform NC^1 , and Extended Frege to $SIZE(poly)$.

In the circuit complexity world, we have lower bounds for explicit functions against DNFs and non-uniform AC^0 , but not against non-uniform NC^1 and $SIZE(poly)$; correspondingly, in the proof complexity setting we have strong lower bounds for Resolution and fairly strong lower bounds for bounded-depth Frege, but no non-trivial lower bounds for Frege and Extended Frege.

The question arises whether this connection between circuit complexity and proof complexity is fundamental or not. No formal connection is known either way — we don’t have theorems to the effect that circuit complexity lower bounds

yield proof complexity lower bounds, nor any implications in the reverse direction. Moreover, barriers such as the natural proofs barrier don't seem to apply to proof complexity. This suggests that perhaps completely different techniques, say from proof theory or finite model theory, might help in showing proof complexity lower bounds. Is there a sense in which there are barriers to making progress in proof complexity? Formulating and understanding such barriers would not only guide us towards the “right” techniques, but might have collateral benefits as well, as in the circuit complexity setting.

In this paper, we shed some light on these questions. We draw a connection between two fundamental lower bound questions in proof complexity. The first question is to prove strong lower bounds for bounded-depth Frege. Superpolynomial lower bounds are known for this proof system, but there aren't any lower bounds known that are purely exponential, i.e., $2^{\Omega(n^c)}$ where the constant c doesn't depend on the depth of lines in the proof (the best known lower bound is $\Omega(2^{n^{5-d}})$ [3]). The second question, which is perhaps the major open question in proof complexity, is to obtain superpolynomial lower bounds for Frege. This question is believed to be very hard — it is non-trivial even to think of plausible candidate tautologies for which superpolynomial lower bounds are believed to hold [4]. We show that progress on the first question would lead to progress on the second, by giving a general simulation of polynomial size Frege proofs by subexponential size bounded-depth Frege proofs. More precisely, we show that even a $2^{n^{\omega(1/d)}}$ proof size lower bound for proving CNF tautologies in depth d Frege would translate to a superpolynomial proof size lower bound for Frege.

The proof of this connection is inspired by a result in circuit complexity, further strengthening the “mapping” between proof complexity and circuit complexity. The circuit complexity result we draw inspiration from is that NC^1 can be simulated by bounded-depth circuits with sub-exponential size [1]. The standard proof of this goes via a divide-and-conquer technique. We use a similar technique in our context, however our task is made harder in a sense by the fact that we need to reason within bounded-depth Frege about equivalence of various alternative representations of a function. The technical heart of our proof involves such reasoning.

Our result is also relevant to algorithmic analysis, which is another major motivation for studying proof complexity. A propositional proof system can be thought of as a non-deterministic algorithm for deciding if a formula is a tautology or not. Proof systems such as bounded-depth Frege and Frege provide particularly simple and natural examples of such algorithms. Indeed, many of the algorithms and heuristics used in practice for solving SAT, such as DPLL and Clause Learning, arise from *determinizing* the non-deterministic algorithm corresponding to some natural proof system. Thus lower bounds for proof systems give us information on the performance of algorithms used in practice.

Algorithmic analysis would appear to be a simpler question than proving complexity lower bounds, since a complexity lower bound is a statement about *any* possible algorithm for a problem, while algorithmic analysis deals with specific algorithms. There are somewhat artificial algorithms such as Levin's optimal

algorithm for SAT whose analysis is just as difficult as proving complexity lower bounds. However, one might expect that for more natural algorithms, such as those corresponding to natural propositional proof systems, this is not the case. Our current lack of progress in proving proof complexity lower bounds indicates that there might be barriers even in algorithmic analysis of natural algorithms. Our main result here can be interpreted as saying that the algorithmic analysis question for the algorithm corresponding to bounded-depth Frege is as hard as the question for the algorithm corresponding to Frege (which in some sense is a more sophisticated algorithm). In general, it would be useful to have a theory of algorithmic analysis which gives us information about the relative difficulty of analyzing various natural algorithms. We make a small step in this direction in the setting of non-deterministic algorithms for TAUT.

There are a couple of interesting byproducts of our main result. First, we are able to prove *tight* bounds for proving certain explicit tautologies in treelike bounded-depth Frege. Lower bounds for the tautologies we consider were already shown by Krajíček [10]. We give corresponding upper bounds as a corollary of our simulation of Frege by bounded-depth Frege.

Second, we address the question of *weak automatizability* for bounded-depth Frege systems. A proof system $\mathcal{P}$ is weakly automatizable if there is an algorithm that on input f and a number r in unary, can distinguish the case where f is not a tautology from the case where f has a $\mathcal{P}$ -proof of size at most r . Despite considerable effort, the question of whether low depth proof systems are weakly automatizable is unresolved. Bonnet, Domingo, Gavalda, Maciel and Pitassi [5] show that depth k Frege systems are not weakly automatizable under a cryptographic assumption, but their result breaks down for small k (less than 6). We use our main result to re-derive their main theorem. Our proof is cleaner and simpler than theirs, and we show that it could potentially resolve the weak automatizability question for lower depth Frege systems than what is currently known.

1.1 Proof Overview

Suppose that P is a Frege proof of some formula f . We want to simulate P by a subexponential-size depth d Frege proof of f . The high-level idea behind the simulation is to replace every formula in the proof by its equivalent depth d (subexponential-size) *flattened* formula, and then to show that if C was derived by a rule from A and B , then the flattened version of C can be efficiently derived from the flattened versions of A and B .

We can assume without loss of generality that all formulas f in the proof are balanced (Reckhow's theorem). We first review the translation of a balanced formula f to its flattened form. We say that a formula has logical depth at most d if the depth of the binary tree representing the formula is at most d . Suppose that we want to replace f , of size n and logical depth $\log n$, by a depth 3 formula. The idea is to view f as consisting of two layers: the top layer is a formula, f_1 , of height $(\log n)/2$, and the bottom layer consists of $2^{(\log n)/2} = \sqrt{n}$ subformulas, $g_1, \dots, g_{\sqrt{n}}$, each of height $(\log n)/2$. Since f_1 has height $(\log n)/2$, it has at most

$\sqrt{n}$ inputs, and thus can be written as either a CNF or a DNF formula (of its inputs) of size $\sqrt{n}2^{\sqrt{n}}$. Similarly, each formula in the bottom layer can be written as either a CNF or a DNF formula of size $\sqrt{n}2^{\sqrt{n}}$. Writing f_1 as a CNF formula, and writing all formulas g_j in the bottom layer as DNF formulas, we obtain a new formula for f of depth 3 and total size $O(n2^{2\sqrt{n}})$. (The depth is 3 because we can merge the middle two AND layers.) In a similar manner, we can replace any formula f , of size n and logical depth $\log n$, by a depth $d + 2$ formula: Now we break up f into d equally-spaced layers, each of size $(\log n)/d$. Again, we write the formula at the top layer as a CNF formula, the formulas at the next layer as DNF formulas, and so on. This gives a formula of depth $2(d + 1)$ and total size $O(n2^{dn^{1/d}})$, but since we alternated CNF/DNFs, we can collapse every other layer to obtain a new flattened formula of depth $2(d + 1) - d = d + 2$.

Now that we have flattened translations of each formula in P , it remains to fill in the proof, to show that the flattened versions can be derived from one another. In order to carry this out, we define a more general procedure for flattening a formula as follows. Let $\mathbf{d}$ be any *depth vector* – i.e., it is a sequence of increasing numbers, where each number in the sequence is between 1 and $\log n$. Then from a balanced formula f of size n and logical depth $\log n$, $\mathbf{d}$ defines a new flattened formula of depth $|\mathbf{d}| + 2$: we break f up into $|\mathbf{d}|$ many levels, where now instead of the levels being equally spaced, the breakpoints are specified by $\mathbf{d}$. For example, if $\mathbf{d} = (4, 12)$ and f has depth 20, then the $\mathbf{d}$ -flattened version of f will have 3 levels, the top level containing levels 1 through 3, the second level 4 through 11, and the third level 12 through 20. Our main lemma shows that for any balanced formula f and any two depth vectors $\mathbf{d}_1, \mathbf{d}_2$, there are efficient low-depth Frege proofs showing that the $\mathbf{d}_1$ -flattened version of f is equivalent to the $\mathbf{d}_2$ -flattened version of f . This main lemma will then allow us to prove that for any rule of our proof system, the flattened versions of the antecedent formulas derive the flattened version of the succedent formula.

2 Proof systems

We will work with the propositional sequent calculus, PK. In the fundamental work of Cook and Reckhow [8], many reasonable formulations of Frege systems (including *all* PK-like systems) were studied and shown to be polynomially equivalent; we work with PK for convenience, but any other Frege system will do.

Each line in a PK proof is a *sequent* of the form $A_1, \dots, A_k \longrightarrow B_1, \dots, B_m$ where $\longrightarrow$ is a new symbol, and A_i, B_j are formulas. The intended meaning is that the conjunction of the A_i 's implies the disjunction of the B_j 's.

A PK *proof* of $\longrightarrow f$ is a sequence of sequents, such that each sequent is either an instance of the axiom $A \longrightarrow A$, or follows from previous sequents from one of the inference rules, and such that the final sequent is $\longrightarrow f$.

The rules of PK are of three types: (i) the structural rules, (ii) the logical rules, and (iii) the cut rule.

The structural rules are weakening, contraction and permutation.

The logical rules allow us to introduce each connective on both the left side and the right side.

The final rule is the cut rule, which allows us to derive $\Gamma \longrightarrow \Delta$ from $A, \Gamma \longrightarrow \Delta$ and $\Gamma \longrightarrow A, \Delta$. We call formula A the *cut formula*.

A full description of PK is found in appendix A.

The *size* of a PK proof is the sum of the sizes of all formulas occurring in the proof.

The *logical depth* of a formula φ , denoted by $\text{ldp}(\varphi)$, is the depth of the formula when considered as a binary tree. For example, $(A \wedge B) \wedge C$ has logical depth 2. A formula whose logical depth is D has size at most $2^{D+1} - 1$, and can depend on at most 2^D variables.

The *depth* of a formula φ , denoted by $\text{dp}(\varphi)$, is the maximum number of alternations between AND and OR connectives from root to leaf, not counting negations, plus one. For example, $(A \wedge B) \wedge C$ has depth 1, and the depth of a CNF or DNF formula is two.

We have given definitions of two different notions of depth. We will use logical depth to reason about formulas in Frege proofs, and depth to reason about formulas in bounded depth proofs.

A *cut-depth k* proof, also called an AC_k^0 -Frege proof, is a PK proof where every *cut formula* in the proof has depth at most k (other formulas are allowed to have arbitrary depth). Note that in the literature, an AC_k^0 -Frege proof is often defined to be a PK proof where *all* formulas have depth at most k . This definition is equivalent to ours if the proven formula has depth at most k .

A PK proof is *tree-like* if the underlying dag structure of the proof forms a tree, i.e. each sequent is used only once.

For technical reasons, we will need all the formulas in our proofs to be balanced. By the following result of Reckhow, this can be assumed without loss of generality.

Theorem 1 (Reckhow, [11, Lemma 4.4.14]). *If a formula of logical depth D has a PK proof of size s , then it has a PK proof of size $s^{O(1)}$ in which all formulas have logical depth $D + O(\log s)$. If the original proof is tree-like, then the new balanced proof is also tree-like.*

Definition 1. *A proof system S is automatizable if there exists an algorithm A such that for all unsatisfiable formulas f , $A(f)$ returns an S -proof of f , and the runtime of A on f is polynomial in the size of the smallest S -proof of f . S is weakly automatizable if there exists a proof system that polynomially simulates S and that is automatizable.*

3 Reducing formula depth

We reduce the depth of a formula using a divide-and-conquer technique. The idea is to decompose the formula into relatively small subtrees, and replace each subtree by a CNF or DNF which is equivalent to the formula computed by the subtree.

Definition 2. Let φ be an arbitrary formula depending on n variables. Denote by $\text{CNF}(\varphi)$ ($\text{DNF}(\varphi)$) some canonically chosen CNF (DNF) representing φ of size $O(n2^n)$. We require that $\text{CNF}(p \wedge q) = \text{DNF}(p \wedge q) = p \wedge q$, and similarly for $p \vee q$ and $\neg p$, when p and q are variables.

We think of formulas as trees in which internal nodes are either binary (if the corresponding connective is $\wedge$ or $\vee$) or unary (when the connective is $\neg$), and leaves are labelled by variables. Each formula has an equivalent formula of the same size where negations only appear immediately above leaves, just by applying De Morgan's laws repeatedly to "move" negations down. We will call such formulas *quasi-monotone*, and will work with them throughout our simulation.

Definition 3. A quasi-monotone formula is one in which negations only appear next to variables, and there are no double negations. Let φ be a quasi-monotone formula. Its dual form $M(\varphi)$ is obtained from φ by switching $\wedge$ and $\vee$ and negating all literals, that is for each variable x switching x and $\neg x$; $M(\varphi)$ is logically equivalent to $\neg\varphi$.

We define two *canonical flattened forms* in parallel.

Definition 4. Let $\mathbf{d} = d_1, \dots, d_k$ be a vector of increasing positive integers. The conjunctive flattened form $C(\varphi; \mathbf{d})$ and disjunctive flattened form $D(\varphi; \mathbf{d})$ of a formula φ are defined recursively as follows. If $k = 0$ (i.e., $\mathbf{d}$ is the empty vector) or $d_1 \geq \text{ldp}(\varphi)$ then $C(\varphi; \mathbf{d}) = \text{CNF}(\varphi)$ and $D(\varphi; \mathbf{d}) = \text{DNF}(\varphi)$. Otherwise, let ψ be the formula obtained from φ by trimming the tree at depth d_1 . The formula ψ depends on the variables of φ as well as on variables corresponding to subformulas of φ at depth d_1 ; we call these true variables and subformula variables, respectively. Let v_χ denote the subformula variable corresponding to the subformula χ .

We explain how to calculate the conjunctive flattened form; the disjunctive flattened form is analogous. Start with $\text{CNF}(\psi)$. Let $\mathbf{e} = d_2 - d_1, \dots, d_k - d_1$. Replace each positive occurrence of a subformula variable v_χ in $\text{CNF}(\psi)$ with $D(\chi; \mathbf{e})$, and each negative occurrence with $M(C(\chi; \mathbf{e}))$. The result is $C(\varphi)$.

The flattened forms are both shallow and not too large.

Definition 5. Let φ be a formula and $\mathbf{d} = d_1, \dots, d_k$ be a vector of increasing positive integers, such that $d_1 \leq \text{ldp}(\varphi)$. Let $d_0 = 0$ and $d_{k+1} = \text{ldp}(\varphi)$. The extent of φ with respect to $\mathbf{d}$ is

$$\text{ex}(\varphi; \mathbf{d}) = \max\{d_{i+1} - d_i : 0 \leq i \leq k\}.$$

Lemma 1. Let φ be a formula and $\mathbf{d}$ a vector of length k and extent $x = \text{ex}(\varphi; \mathbf{d})$. Then $C(\varphi; \mathbf{d})$ and $D(\varphi; \mathbf{d})$ are formulas of depth at most $k + 2$ and size $2^{O(k2^x)}$ equivalent to φ .

4 Proof of Main Theorem

In this section, we will prove the following theorem.

Theorem 2. *Let φ be a formula provable in Frege in size s , satisfying $\text{ldp}(\varphi) \leq C \log s$. For every $k \geq 1$ there is an AC_{k+2}^0 -Frege proof of φ of size $2^{O(k s^{O(C/k)})}$. Furthermore, if the original proof is tree-like, so is the new one.*

Corollary 1. *Let φ be a formula of size s and logical depth at most $C \log s$. If φ has a Frege proof of size $O(s^c)$ then for every $k \geq 1$ there is an AC_{k+2}^0 -Frege proof of φ of size $2^{O(ck s^{O(C/k)})}$.*

We will first state some simple lemmas which will enable us to reason about flattened forms. The proofs of these lemmas appear in the Appendix.

Lemma 2. *Let $\Gamma \longrightarrow \Delta$ be a valid sequent of size m , in which n variables appear. The sequent is provable using a tree-like proof of size $O(m^2 n 2^n)$ which cuts only on variables.*

Our next lemma states that we can substitute formulas for variables to get a valid proof.

Lemma 3. *Let π be a proof of $\Gamma \longrightarrow \Delta$ of size s , and let x be a variable appearing in $\Gamma \longrightarrow \Delta$. If we substitute everywhere a formula φ of size m for x then we get a valid proof of size at most sm .*

The preceding lemma shows that we can lift a proof of a sequent by attaching stuff ‘below’. The next lemma shows that we can also lift a proof by attaching stuff ‘above’; this corresponds to deep inference.

Definition 6. *The double sequent $P \longleftrightarrow Q$ is the pair of sequents $P \longrightarrow Q$ and $Q \longrightarrow P$.*

Lemma 4. *Let $P \longrightarrow Q$ be a sequent of size m , and $\varphi(x)$ be a formula of size n in which the variable x appears only once (other variables may also appear). The double sequent $\varphi(x|P) \longleftrightarrow \varphi(x|Q)$ has a cut-free, tree-like proof from the double sequent $P \longleftrightarrow Q$ of size $O(n(m+n))$ (this means that each of $P \longrightarrow Q$ and $Q \longrightarrow P$ is used only once in the joint proof).*

We next state two easy lemmas on dualization.

Lemma 5. *Let φ be a quasi-monotone formula of size n . The double sequent $M(\varphi) \longleftrightarrow \neg \varphi$ has a cut-free, tree-like proof of size $O(n^2)$.*

The second lemma allows us to lift an equivalence to its dualized version.

Lemma 6. *Let φ, ψ be quasi-monotone formulas. Suppose that the double sequent $\varphi \longleftrightarrow \psi$ has a proof of size s cutting on formulas of depth at most D . Then the double sequent $M(\varphi) \longleftrightarrow M(\psi)$ has a proof of size $O(s)$ cutting on formulas of depth at most D . Furthermore, if the original proof is tree-like then so is the new proof.*

We comment that the preceding lemma can be strengthened to produce cut-free proofs.

4.1 Moving down the depth vector

In this section we show how to prove the equivalence of two flattened forms of the same formula which correspond to two different depth vectors.

Lemma 7. *Let φ be a formula of logical depth D , and δ a positive integer. Consider $\text{CNF}(\varphi)$ and $\text{C}(\varphi; \delta)$ as monotone formulas depending on literals $x, \bar{x}$; in other words, for each variable x , we replace $\neg x$ by $\bar{x}$. The double sequent $\text{CNF}(\varphi) \longleftrightarrow \text{C}(\varphi; \delta)$ has a tree-like proof of size $2^{O(2^D)}$ cutting only on literals.*

Proof. Since φ has at most 2^D leaves, it depends on at most 2^D variables, and twice as many literals. Lemma 2 provides the necessary proof of size $2^{O(D)} 2^{2^{D+1}} = 2^{O(2^D)}$.

Lemma 8. *Let φ be a formula, $\mathbf{d} = d_1, \dots, d_k$ a vector of increasing positive integers, and $\delta < d_1$ be a positive integer. The double sequent $\text{C}(\varphi; \mathbf{d}) \longleftrightarrow \text{C}(\varphi; \delta, \mathbf{d})$ has a tree-like proof of size $2^{O(k2^x)}$ cutting only on formulas of depth at most $k+1$, where $x = \text{ex}(\varphi; \mathbf{d})$.*

Proof. Let ψ be φ trimmed at level d_1 , as in Definition 4. By definition, $\text{C}(\varphi; \mathbf{d})$ is obtained from $\text{CNF}(\psi)$ by substituting for each subformula literal $v, \bar{v}$ a formula v^+, v^- of depth $k+1$. Similarly, $\text{C}(\varphi; \delta, \mathbf{d})$ is obtained by the same substitution to $\text{C}(\psi; \delta)$.

Lemma 7 shows how to prove $\text{CNF}(\psi) \longleftrightarrow \text{C}(\psi; \delta)$ cutting only on literals; note that $\text{ldp}(\psi) \leq x$. By Lemma 3, if we substitute v^+, v^- for $v, \bar{v}$ we obtain a valid proof of size $2^{O(k2^x)}$. The cut formulas lift from literals to v^+, v^- , which are of depth at most $k+1$.

Lemma 9. *Let φ be a formula, $\mathbf{d} = d_1, \dots, d_k$ a vector of increasing positive integers, and $d_i < \delta < d_{i+1}$, where $1 \leq i \leq k$. Define $\mathbf{e} = d_1, \dots, d_i, \delta, d_{i+1}, \dots, d_k$. The double sequent has a tree-like proof of size $2^{O(k2^x)}$ cutting only on formulas of depth at most $k+1$, where $x = \text{ex}(\varphi; \mathbf{d})$.*

Proof. Let ψ be the portion of φ up to level d_i . Both flattened forms in consideration are obtained from $\text{C}(\psi; d_1, \dots, d_i)$ by substituting certain formulas for literals; these formulas are either flattened forms or dualized flattened forms. Denote the formulas corresponding to the literals $v, \bar{v}$ by v^+, v^- .

We can assume wlog that v^+, v^- are CNFs. In order to obtain $\text{C}(\varphi; \mathbf{d})$, we need to substitute $v_d^+ = \text{CNF}(v; d_{i+1}, \dots, d_k)$. In order to obtain $\text{C}(\varphi; \mathbf{e})$, we need to substitute $v_e^+ = \text{CNF}(v; \delta, d_{i+1}, \dots, d_k)$. By Lemma 8, we can prove $v_d^+ \longleftrightarrow v_e^+$ in size $2^{O(k2^x)}$, cutting only on formulas of depth at most $k+1$.

Next, $v_d^- = \text{M}(\text{DNF}(v; d_{i+1}, \dots, d_k))$ and $v_e^- = \text{M}(\text{DNF}(v; \delta, d_{i+1}, \dots, d_k))$. Combining Lemma 8 with Lemma 6, we obtain a proof of $v_d^- \longleftrightarrow v_e^-$ in size $2^{O(k2^x)}$. Since the depth of formulas we cut on is preserved in the proof of Lemma 6, the new proof has cuts only on formulas of depth at most $k+1$.

Define now hybrid formulas χ_t as follows. Start with $\text{C}(\psi; d_1, \dots, d_i)$, and replace the first t occurrences of subformula literals by the corresponding $v_e^\pm$;

replace the rest by the corresponding $v_d^\pm$. Thus $\chi_0 = C(\varphi; \mathbf{d})$, and for some $T \leq 2^{O(k2^x)}$, $\chi_T = C(\varphi; \mathbf{e})$.

Using Lemma 4, for any $0 \leq t < T$ we can prove $\chi_t \longleftrightarrow \chi_{t+1}$ given an instance of the corresponding double sequent $v_d^\pm \longleftrightarrow v_e^\pm$. By cutting on all χ_t for $0 < t < T$, we obtain a proof of the desired double sequent.

Lemma 10. *Let φ be a formula, and $\mathbf{d} = d_1, \dots, d_k$ be a vector of increasing positive integers. Define $\mathbf{e} = 1, d_1 + 1, \dots, d_k + 1$. The double sequent $C(\varphi; \mathbf{d}) \longleftrightarrow C(\varphi; \mathbf{e})$ has a tree-like proof of size $2^{O(k2^x)}$ cutting only on formulas of depth at most $k + 3$, where $x = \max(\text{ex}(\varphi; \mathbf{d}), \text{ex}(\varphi; \mathbf{e}))$.*

Proof. Lemma 9 shows how to prove the equivalence of two flattened forms, where the second one has one extra level beyond the first one. In order to ‘move’ d_1 to $d_1 + 1$, we prove the following double sequents:

$$\begin{aligned} C(\varphi; d_1, \dots, d_k) &\longleftrightarrow C(\varphi; d_1, d_1 + 1, \dots, d_k), \\ C(\varphi; d_1, d_1 + 1, \dots, d_k) &\longleftrightarrow C(\varphi; d_1 + 1, \dots, d_k). \end{aligned}$$

Continuing the same way, we can ‘migrate’ $\mathbf{d}$ to $\mathbf{e}$, adding the extra $e_1 = 1$ at the end. Each flattened form in the interim has depth at most $k + 3$. By cutting all the intermediate flattened forms, we obtain the desired double sequent.

The same methods used to prove Lemma 9 enable us to prove the following lemma.

Lemma 11. *Let φ be a formula, and $\mathbf{d} = d_1, \dots, d_k$ be a vector of increasing positive integers. The double sequent $C(\varphi; \mathbf{d}) \longleftrightarrow D(\varphi; \mathbf{d})$ has a tree-like proof of size $2^{O(k2^x)}$ cutting only on formulas of depth at most $k + 2$, where $x = \text{ex}(\varphi; \mathbf{d})$.*

4.2 Putting it together

In this section we show how to transform a Frege proof to an AC^0 -Frege proof. We begin by proving intensional comprehension.

Lemma 12. *Let φ, ψ be formulas, and $\mathbf{d}$ be a vector of increasing positive integers of length k . The double sequents*

$$C(\varphi \wedge \psi; \mathbf{d}) \longleftrightarrow C(\varphi; \mathbf{d}) \wedge C(\psi; \mathbf{d}), \quad C(\varphi \vee \psi; \mathbf{d}) \longleftrightarrow C(\varphi; \mathbf{d}) \vee C(\psi; \mathbf{d})$$

have tree-like proofs of size $2^{O(k2^x)}$ with cuts on formulas of depth at most $k + 3$, where $x = \text{ex}(\varphi \wedge \psi; \mathbf{d})$.

Proof. We show how to prove the first double sequent; the second is proven in the same way.

Let $\mathbf{e} = 1, d_1 + 1, \dots, d_k + 1$ be the vector defined in Lemma 10. We calculate $D(\varphi \wedge \psi; \mathbf{e})$. Using the recipe of definition 4, we first calculate $\text{DNF}(v_\varphi \wedge v_\psi) = v_\varphi \wedge v_\psi$. Into this DNF we substitute $v_\varphi = C(\varphi; \mathbf{d})$ and $v_\psi = C(\psi; \mathbf{d})$. Therefore $D(\varphi \wedge \psi; \mathbf{e}) = C(\varphi; \mathbf{d}) \wedge C(\psi; \mathbf{d})$. The proof now becomes obvious, along the following lines.

Using Lemma 11, we prove $C(\varphi \wedge \psi; \mathbf{d}) \longleftrightarrow D(\varphi \wedge \psi; \mathbf{d})$. Using Lemma 10, we prove $D(\varphi \wedge \psi; \mathbf{d}) \longleftrightarrow C(\varphi; \mathbf{d}) \wedge C(\psi; \mathbf{d})$. The proof is completed by an application of the cut rule.

Lemma 13. *Let φ be a formula, and $\mathbf{d}$ be a vector of increasing positive integers of length k . The double sequent $C(\neg\varphi; \mathbf{d}) \longleftrightarrow \neg C(\varphi; \mathbf{d})$ has a tree-like proof of size $2^{O(k2^x)}$ with cuts on formulas of depth at most $k + 3$, where $x = \text{ex}(\neg\varphi; \mathbf{d})$.*

Proof. Let $\mathbf{e} = 1, d_1 + 1, \dots, d_k + 1$ be the vector defined in Lemma 10. We calculate $C(\neg\varphi; \mathbf{e})$. Using the recipe of definition 4, we first calculate $\text{CNF}(\neg v_\varphi) = \neg v_\varphi$. Into this CNF we substitute $\neg v_\varphi = M(C(\varphi; \mathbf{d}))$. Therefore $C(\neg\varphi; \mathbf{e}) = M(C(\varphi; \mathbf{d}))$. The proof now becomes obvious, along the following lines.

Lemma 10 shows how to prove $C(\neg\varphi; \mathbf{d}) \longleftrightarrow M(C(\varphi; \mathbf{d}))$. Lemma 5 shows how to prove $M(C(\varphi; \mathbf{d})) \longleftrightarrow \neg C(\varphi; \mathbf{d})$. The proof is completed by applying the cut rule.

The preceding lemmas allow us to unroll flattened forms.

Lemma 14. *Let φ be a formula, and $\mathbf{d}$ be a vector of increasing positive integers of length k . The double sequent $\varphi \longleftrightarrow C(\varphi; \mathbf{d})$ has a tree-like proof of size $2^{O(k2^x)}$ with cuts on formulas of depth at most $k + 3$, where $x = \text{ex}(\varphi; \mathbf{d})$.*

Proof. The proof is by structural induction. If φ is a literal then there is nothing to prove. If $\varphi = \neg\psi$, then use Lemma 13 to prove $C(\varphi; \mathbf{d}) \longleftrightarrow \neg C(\psi; \mathbf{d})$. The induction hypothesis gives us a proof of $\psi \longleftrightarrow C(\psi; \mathbf{d})$; move both ψ and its flattened form to the other side using four $\neg$ introduction rules, and apply cut twice to prove the required double sequent.

If $\varphi = \psi \wedge \chi$ then start with proofs of the following sequents, obtained by Lemma 12 and the induction hypothesis:

$$C(\varphi; \mathbf{d}) \longleftrightarrow C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d}), \quad \psi \longleftrightarrow C(\psi; \mathbf{d}), \quad \chi \longleftrightarrow C(\chi; \mathbf{d}).$$

Now prove $C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d}) \longleftrightarrow \psi \wedge \chi$ as follows:

$$\frac{\frac{C(\psi; \mathbf{d}) \longrightarrow \psi}{C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d}) \longrightarrow \psi} \wedge L \quad \frac{C(\chi; \mathbf{d}) \longrightarrow \chi}{C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d}) \longrightarrow \chi} \wedge L}{C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d}) \longrightarrow \psi \wedge \chi} \wedge R$$

The other sequent is proved similarly. Complete the proof using the cut rule. The case $\varphi = \psi \vee \chi$ is similar.

The proof of the main theorem is now simple.

Lemma 15. *Let φ be a formula provable in Frege in size s using a proof with maximum logical depth D . For every k there is an AC_{k+2}^0 -Frege proof of φ of size $s2^{O(k2^{D/k})}$. Furthermore, if the original proof is tree-like, so is the new one.*

Proof. Let $\mathbf{d} = \lceil D/k \rceil, 2\lceil D/k \rceil, \dots, (k-1)\lceil D/k \rceil$. Note that the extent of each formula with respect to $\mathbf{d}$ is at most $x = \lceil D/k \rceil$. Take the original proof and replace each formula ψ by $C(\psi; \mathbf{d})$. Each application of a rule is still valid, but the proof as a whole isn't valid since not all formulas are in flattened form. We address this issue by tampering with the introduction rules, as in the following example, corresponding to the right $\wedge$ introduction rule:

$$\frac{\Gamma \longrightarrow \Delta, C(\psi; \mathbf{d}) \quad \Gamma \longrightarrow \Delta, C(\chi; \mathbf{d})}{\Gamma \longrightarrow \Delta, C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d})} \wedge R \quad \frac{\Gamma \longrightarrow \Delta, C(\psi; \mathbf{d}) \wedge C(\chi; \mathbf{d})}{\Gamma \longrightarrow \Delta, C(\psi \wedge \chi; \mathbf{d})} \text{Lem. 12 Cut}$$

Applying the same transformation for all introduction rules, we are left with a valid proof of $\longrightarrow C(\varphi; \mathbf{d})$, where each sequent is now replaced by sequents of total size $2^{O(k2^x)}$; the total size so far is $s2^{O(k2^x)}$. Lemma 14 proves $C(\varphi; \mathbf{d}) \longrightarrow \varphi$, and the proof is complete by cutting on $C(\varphi; \mathbf{d})$.

The lemmas we used employ cuts of depth at most $k+2$. All cuts in the original proof now cut flattened formulas, which are of depth at most $k+1$.

Proof (of Theorem 2). Reckhow's Theorem (Theorem 1) supplies us with an $AC_{O(\log s)}^0$ proof of φ of size $s^{O(1)}$. The theorem now follows by substituting $D = C \log(s)$ in Lemma 15.

5 Applications and Consequences

5.1 Tightness of our simulation

We first address the tightness of our simulation. The analogous result for circuit complexity shows that any function computable by a polynomial-size formula can be computed by depth d circuits of size $\exp(n^{O(1/d)})$. This result is tight, since Håstad's theorem proves that the parity function on n boolean variables requires AC_d^0 circuits of size $\exp(n^{1/d})$.

Similarly we can show that our result is also tight. The following theorem states that there are formulas that have polynomial-size Frege proofs, but that require AC_d^0 proofs of size exponential in $n^{1/d}$.

Theorem 3. *For every d there is a sequence of balanced formulas φ_n of depth $d+2$ provable in Frege by a tree-like proof of size s_n such that every tree-like AC_d^0 proof of φ_n requires size $2^{s_n^{\Omega(1/d)}}$.*

Proof. The formula φ_n is PHP_n , the pigeonhole principle with $n+1$ pigeons and n holes, with each variable replaced by a Sipser function of depth d . Buss [7] has shown how to prove PHP_n using a Frege proof of size $n^{O(1)}$, which can be made tree-like by squaring its size. Substituting the Sipser functions, we obtain a Frege proof of size $n^{d+O(1)}$.

Conversely, Krajíček [10] gives a lower bound of $2^{n^{\Omega(1)}}$ for proving φ_n in tree-like AC_d^0 .

Since the formulas φ_n are balanced, Theorem 2 applies, and with $k = d - 2$, gives proofs essentially matching the lower bound.

The above result proves tightness for formulas of high depth. We conjecture that our simulation is also tight with respect to CNF formulas and general, dag-like proofs. The obvious formula for witnessing the lower bound is the pigeonhole principle itself. However, as an artifact of the switching lemma technique used to obtain depth d Frege lower bounds for the pigeonhole principle, the current best lower bound is exponential in $n^{1/2^d}$. It is a well-known open problem to improve the lower bound to $\exp(n^{1/d})$ for the pigeonhole principle, or for any other CNF formula. Such a result would show that our simulation is tight even for CNF formulas and arbitrary dag-like proofs.

5.2 Weak Automatizability

Using our theorem, we are able to show that bounded-depth Frege is not weakly automatizable, under an assumption about the hardness of factoring. While this result has already been known [5], we first show how to prove it as a simple corollary of our main theorem.

Theorem 4 ([6]). *Frege systems do not have feasible interpolation and are not weakly automatizable unless the Diffie Hellman problem is computable by polynomial size circuits.*

The Diffie Hellman problem is based on a prime number p , $|p| = n$. The input to the problem is a number g less than p , and numbers $g^a \pmod{p}$, $g^b \pmod{p}$, for some numbers $a, b \leq p$. The output should be $g^{ab} \pmod{p}$. The main lemma from [6] shows that a particular tautology, DH_p , stating that the Diffie Hellman function is well defined, has Frege proofs of size $O(|p|^c)$, where $c \leq 4$. Take DH_p where $|p| = (\log n)^q$ for some constant q . By our normal form theorem, this implies that DH_p has AC_k^0 -Frege proofs of size $2^{O(k(\log n)^{cq/k})}$. Thus for $k > cq$, this is polynomial in n . Hence it follows that if AC_k^0 -Frege is weakly automatizable (or has feasible interpolation), then the Diffie Hellman problem for $|p| = n' = (\log n)^{k/c}$ can be solved in time $n^{O(k)} = 2^{O(k \log n)} = \exp O(k(n')^{c/k})$.

Unfortunately, the quality of this negative result degrades for small k . Indeed despite considerable effort, it is unknown whether or not very low depth Frege systems (when k is less than 5) are weakly automatizable (the recent paper [2] reveals a connection between AC_k^0 with bottom fan-in 2 and mean-payoff games). The main reason for this is that the Diffie Hellman function is not hard enough! Algorithms exist for computing discrete log over all finite fields, and hence for Diffie Hellman, that run in time exponential in $\sqrt{n}$. Moreover, the number field sieve is conjectured to solve discrete log (and thus Diffie Hellman) in time exponential in the cube root of n . On the other hand, it seems entirely possible to come up with a different interpolant statement for another function that is much harder – truly exponential in n , and that still has efficient Frege proofs. Using our main theorem (which scales down *any* Frege proof), this would imply new negative results for weak automatizability and feasible interpolation for lower depth Frege systems than what is currently known.

References

1. Eric Allender, Lisa Hellerstein, Paul McCabe, Toniann Pitassi, and Michael Saks. Minimizing disjunctive normal form formulas and AC^0 circuits given a truth table. *SIAM Journal on Computing*, 38(1):63–84, 2008.
2. Albert Atserias and Elitza Maneva. Mean-payoff games and propositional proofs. *Inf. and Comp.*, 209(4):664–691, 2011.
3. Paul W. Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, Pavel Pudlák, and Alan Woods. Exponential lower bounds for the pigeonhole principle. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Computing*, pages 200–220, Victoria, B.C., Canada, May 1992.
4. Maria Luisa Bonet, Samuel R. Buss, and Toniann Pitassi. Are there hard examples for Frege systems? In *Feasible Mathematics II*, pages 30–56. Birkhäuser, 1995.
5. Maria Luisa Bonet, Carlos Domingo, Ricard Gavaldà, Alexis Maciel, and Toniann Pitassi. Non-automatizability of bounded-depth Frege proofs. *Computational Complexity*, 13(1-2):47–68, 2004.
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.
7. Samuel R. Buss. Polynomial size proofs of the pigeonhole principle. *Journal of Symbolic Logic*, 57:916–927, 1987.
8. Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, 1979.
9. Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–305, 1985.
10. Jan Krajíček. Lower bounds to the size of constant-depth propositional proofs. *Journal of Symbolic Logic*, 59(1):73–86, March 1994.
11. Jan Krajíček. *Bounded arithmetic, propositional logic, and complexity theory*. Cambridge University Press, New York, NY, USA, 1995.

A The Proof System PK

The *formulas* in PK are formed from variables, the binary connectives $\wedge$ and $\vee$ and the unary connective $\neg$. Each line in a PK proof is a *sequent* of the form $A_1, \dots, A_k \longrightarrow B_1, \dots, B_m$ where $\longrightarrow$ is a new symbol, and A_i, B_j are formulas. The intended meaning is that the conjunction of the A_i ’s implies the disjunction of the B_j ’s. Thus, a proof of f in PK will be interpreted to be a PK proof of the sequent $\longrightarrow f$.

A PK *proof* of $\longrightarrow f$ is a sequence of sequents, such that each sequent is either an instance of the axiom $A \longrightarrow A$, or follows from previous sequents from one of the inference rules, and such that the final sequent is $\longrightarrow f$.

The rules of PK are of three types: (i) the structural rules, (ii) the logical rules, and (iii) the cut rule.

The structural rules are weakening (formulas can always be added to the left or to the right), contraction (two copies of the same formula can be replaced by one) and permutation (formulas in a sequent can be reordered).

The final rule is the cut rule, which allows us to derive $\Gamma \longrightarrow \Delta$ from $A, \Gamma \longrightarrow \Delta$ and $\Gamma \longrightarrow A, \Delta$. The formula A is called the *cut formula*.

The logical rules, shown below, allow us to introduce each connective on both the left side and the right side.

1. (Negation Left, $\neg\mathbf{L}$) From $\Gamma \longrightarrow A, \Delta$, derive $\neg A, \Gamma \longrightarrow \Delta$.
2. (Negation Right, $\neg\mathbf{R}$) From $A, \Gamma \longrightarrow \Delta$, derive $\Gamma \longrightarrow \neg A, \Delta$.
3. (And Left, $\wedge\mathbf{L}$) From $A, B, \Gamma \longrightarrow \Delta$, derive $A \wedge B, \Gamma \longrightarrow \Delta$.
4. (And Right, $\wedge\mathbf{R}$) From $\Gamma \longrightarrow A, \Delta$ and $\Gamma \longrightarrow B, \Delta$, derive $\Gamma \longrightarrow A \wedge B, \Delta$.
5. (Or Left, $\vee\mathbf{L}$) From $A, \Gamma \longrightarrow \Delta$ and $B, \Gamma \longrightarrow \Delta$, derive $A \vee B, \Gamma \longrightarrow \Delta$.
6. (Or Right, $\vee\mathbf{R}$) From $\Gamma \longrightarrow A, B, \Delta$ derive $\Gamma \longrightarrow A \vee B, \Delta$.

The *size* of a PK proof is the sum of the sizes of all formulas occurring in the proof. The *logical depth* of a formula φ , denoted by $\text{ldp}(\varphi)$, is the depth of the formula when considered as a tree. For example, $(A \wedge B) \wedge C$ has logical depth 2. A formula whose logical depth is D has size at most $2^{D+1} - 1$, and can depend on at most 2^D variables. The *depth* of a formula φ , denoted by $\text{dp}(\varphi)$, is the maximum number of alternations from root to leaf, not counting negations, plus one. For example, $(A \wedge B) \wedge C$ has depth 1, and the depth of a CNF or DNF formula is two. We have given two different, incompatible definitions of depth. We will use logical depth to reason about formulas in Frege proofs, and depth to reason about formulas in bounded depth proofs.

We will think of a formula of depth k as an unbounded fan-in formula. A formula is Σ_k if it has depth k , where the top (unbounded fan-in) connective is OR; a formula is Π_k if it has depth k and the top (unbounded fan-in) connective is AND. A formula is Δ_k if it can be written both as a Σ_k and as a Π_k formula.

A depth k proof is a PK proof where every formula in the proof has depth at most k . Such a proof is usually called an AC_k^0 -Frege proof, but for us this term will have a slightly different meaning. Similarly we can define Σ_k , Π_k , and Δ_k proofs.

In the above definition of depth k proofs, *all* formulas in the proof are required to have depth at most k . A more general definition of depth can also be defined in order to allow constant depth proofs of higher depth formulas. Let f be a formula, of possibly high depth. A *cut-depth k* Frege proof of f , also known as an AC_k^0 -Frege proof, is a PK proof of $\longrightarrow f$ such that in all applications of the cut rule, the cut formula has depth at most k . Similarly, we can define proofs of cut-depth Σ_k , Π_k and Δ_k . This new definition generalizes our earlier definition of depth k proofs because they can be shown to be equivalent when restricted to formulas f of depth k .

Underlying a PK proof is a directed acyclic graph representing the implication structure of the proof. (Each node in the graph corresponds to a sequent in the proof, and edges $(s_1, s_3), (s_2, s_3)$ means that the sequent s_3 is derived from sequents s_1 and s_2 via an application of one of the PK rules.) A PK proof is *tree-like* if the underlying dag structure of the proof forms a tree. When presenting proofs in PK, we will only mention the logical rules and the cut rule, but not the structural rules.

B Proofs of Lemmas from Section 3

Lemma 1. Let φ be a formula and $\mathbf{d}$ a vector of length k and extent $x = \text{ex}(\varphi; \mathbf{d})$. Then $C(\varphi; \mathbf{d})$ and $D(\varphi; \mathbf{d})$ are formulas of depth at most $k + 2$ and size $2^{O(k2^x)}$ equivalent to φ .

Proof. It is easy to see, using De Morgan's laws, that the flattened forms are equivalent to the original formula. The recursive definition of the flattened forms ensures that all negations are pushed to the leaves, and that CNFs and DNFs alternate. Therefore their depth is $k + 2$ (the depth of a CNF/DNF is 2).

In order to estimate the size, denote by $M(k, x)$ the maximum size of a flattened form of a formula with respect to a vector of length k and extent x . By definition 2,

$$M(0, x) = O(2^x 2^{2^x}) = 2^{O(2^x)},$$

since a formula of logical depth x depends on at most 2^x variables. Since $M(0, x)$ also bounds the number of literals in a CNF/DNF, $M(k + 1, x) \leq M(0, x) + M(0, x)M(k, x)$. Therefore

$$M(k, x) \leq \sum_{l=0}^k M(0, x)^{l+1} = 2^{O(k2^x)}. \square$$

C Proofs of Lemmas from Section 4

In order to prove the lemmas from this section, we will need a few preliminary lemmas.

Definition 7. A truth assignment for variables $x_1, \dots, x_n$ is a function $f: \{x_1, \dots, x_n\} \rightarrow \{\perp, \top\}$ assigning to each variable a truth value ($\perp$ is False, $\top$ is True).

Lemma 16. Let φ be a formula of size m depending upon the set of variables X of size n , and consider a truth assignment f for X . If φ is satisfied by f then the sequent

$$\{x \in X : f(x) = \perp\} \rightarrow \{x \in X : f(x) = \top\}, \varphi$$

has a cut-free, tree-like proof of size $O(nm(n + m))$. If φ is falsified by f , then the same is true for the sequent

$$\{x \in X : f(x) = \perp\}, \varphi \rightarrow \{x \in X : f(x) = \top\}.$$

Proof. The proof is by structural induction. Denote by $S(\varphi)$ the sequent alluded to in the statement of the lemma. We first describe the proof, and then analyze its size.

If $\varphi = x$ is a variable then $S(\varphi)$ follows from the axiom $x \rightarrow x$ using structural rules. If $\varphi = \neg\psi$ then $S(\varphi)$ follows from $S(\psi)$ by using the appropriate $\neg$ introduction rule.

If $\varphi = \psi \wedge \chi$ and φ is satisfied by f , then $S(\varphi)$ follows from the sequents $S(\psi)$ and $S(\chi)$ by using the right $\wedge$ introduction rule. If it is falsified by f , then either ψ is falsified or χ is falsified. Suppose wlog that ψ is falsified. Then $S(\varphi)$ follows from $S(\psi)$ by using the left $\wedge$ introduction rule. The proofs are similar if the main connective is $\vee$ instead of $\wedge$. In total, we have eliminated each connective by using one logical rule, and each variable using n weakening rules. The total number of sequents needed is therefore $O(nm)$, each of size at most $n + m$.

Lemma 17. *Let $\Gamma \longrightarrow \Delta$ be a sequent of size m depending upon the set of variables X of size n . Suppose that the sequent is valid under some truth assignment f . Then the sequent*

$$\Gamma, \{x \in X : f(x) = \perp\} \longrightarrow \Delta, \{x \in X : f(x) = \top\}$$

has a cut-free, tree-like proof of size $O(nm(n + m))$.

Proof. Since the sequent is valid under f , either one of the formulas in Γ is false, or one of the formulas in Δ is true. Use Lemma 16 to prove the corresponding sequent, and conclude the sequent in the statement by using at most m weakening rules, for an extra size of $O(m(n + m))$.

Lemma 2. Let $\Gamma \longrightarrow \Delta$ be a valid sequent of size m , in which n variables appear. The sequent is provable using a tree-like proof of size $O(m^2 n 2^n)$ which cuts only on variables.

Proof. Let X be the set of variables appearing in $\Gamma \longrightarrow \Delta$; note that $n = |X| \leq m$. Apply Lemma 17 for each of the 2^n truth assignments. Divide all truth assignments into pairs where only the value of the leftmost variable in $x \in X$ differs. Apply the cut rule to all pairs, eliminating the variable x . Continue this way, eliminating all variables in order, to obtain a proof of $\Gamma \longrightarrow \Delta$. In total, the proof uses $2^n - 1$ cuts.

Lemma 3. Let π be a proof of $\Gamma \longrightarrow \Delta$ of size s , and let x be a variable appearing in $\Gamma \longrightarrow \Delta$. If we substitute everywhere a formula φ of size m for x then we get a valid proof of size at most sm .

Proof. All the rules of PK are closed under substitution.

Lemma 4. Let $P \longrightarrow Q$ be a sequent of size m , and $\varphi(x)$ be a formula of size n in which the variable x appears only once (other variables may also appear). The double sequent

$$\varphi(x|P) \longleftrightarrow \varphi(x|Q)$$

has a cut-free, tree-like proof from the double sequent $P \longleftrightarrow Q$ of size $O(n(m + n))$ (this means that each of $P \longrightarrow Q$ and $Q \longrightarrow P$ is used only once in the joint proof).

Proof. The proof is by structural induction. If $\varphi = x$ then there is nothing to prove. If $\varphi = \neg\psi$, then $\varphi(P) \longleftrightarrow \varphi(Q)$ follows from $\psi(P) \longleftrightarrow \psi(Q)$ by four applications of the $\neg$ introduction rules.

If $\varphi = \psi \wedge \chi$, then assume wlog that x appears in χ . We use the following proof twice:

$$\frac{\frac{\psi \longrightarrow \psi}{\psi \wedge \chi(P) \longrightarrow \psi} \wedge L \quad \frac{\frac{\chi(P) \longrightarrow \chi(Q)}{\psi \wedge \chi(P) \longrightarrow \chi(Q)} \wedge L}{\psi \wedge \chi(P) \longrightarrow \psi \wedge \chi(Q)} \wedge R$$

Each instance of the proof uses a different assumption. A similar proof works if the main connective is $\vee$ instead of $\wedge$. In total, there are $O(n)$ sequents of size $O(m+n)$.

Lemma 5. Let φ be a quasi-monotone formula of size n . The double sequent

$$M(\varphi) \longleftrightarrow \neg\varphi$$

has a cut-free, tree-like proof of size $O(n^2)$.

Proof. We construct inductively proofs of the double sequent $M(\varphi), \varphi \longleftrightarrow$. From this, we conclude the double sequent $M(\varphi) \longleftrightarrow \neg\varphi$ using two applications of the $\neg$ introduction rules.

If $\varphi = x$ or $\varphi = \neg x$ then required double sequent is proved as follows:

$$\frac{x \longrightarrow x}{\neg x, x \longrightarrow} \neg L \quad \frac{x \longrightarrow x}{\longrightarrow \neg x, x} \neg R.$$

If $\varphi = \psi \wedge \chi$ then the proof is

$$\frac{\frac{\frac{M(\psi), \psi \longrightarrow}{M(\psi), \psi \wedge \chi \longrightarrow} \wedge L \quad \frac{\frac{M(\chi), \chi \longrightarrow}{M(\chi), \psi \wedge \chi \longrightarrow} \wedge L}{M(\psi) \vee M(\chi), \psi \wedge \chi \longrightarrow} \vee L}{\longrightarrow M(\psi), \psi} \vee R \quad \frac{\longrightarrow M(\chi), \chi}{\longrightarrow M(\psi) \vee M(\chi), \chi} \vee R}{\longrightarrow M(\psi) \vee M(\chi), \psi \wedge \chi} \wedge R$$

If $\varphi = \psi \vee \chi$ then the proof is similar. In all, we have $O(n)$ sequents of size $O(n)$.

We comment that the preceding lemma can be strengthened to produce cut-free proofs, essentially by replacing each instance of the (negated) axiom $\neg x \longrightarrow \neg x$ with $\neg x, x \longrightarrow \neg x, x$.

Lemma 6. Let φ, ψ be quasi-monotone formulas. Suppose that the double sequent $\varphi \longleftrightarrow \psi$ has a proof of size s cutting on formulas of depth at most D . Then the double sequent $M(\varphi) \longleftrightarrow M(\psi)$ has a proof of size $O(s)$ cutting on formulas of depth at most D ; the new proof can include cuts even if the original one was cut-free. Furthermore, if the original proof is tree-like then so is the new proof.

Proof. Denote by $\bar{\varphi}, \bar{\psi}$ the formulas obtained from φ, ψ by negating all literals. If we negate all literals in the proof of $\varphi \longleftrightarrow \psi$ then the proof almost remains valid, now proving $\bar{\varphi} \longleftrightarrow \bar{\psi}$; the only problematic rules are the negation introduction rules, when applied to variables. If the original rule replaced a variable x by its negation $\neg x$, then the new rule is supposed to replace $\neg x$ with x . This is realized by a derivation which replace $\neg x$ with $\neg\neg x$, which is then replaced by x through cutting with $\neg\neg x \longrightarrow x$. Multiple negations are handled similarly.

If we take the proof of $\bar{\varphi} \longrightarrow \bar{\psi}$, switch $\wedge$ with $\vee$, and switch the side of each formula in the proof, then we get a valid proof of $M(\psi) \longrightarrow M(\varphi)$. Similarly we can obtain a proof of the other sequent.

Lemma 7. Let φ be a formula of logical depth D , and δ a positive integer. Consider $\text{CNF}(\varphi)$ and $C(\varphi; \delta)$ as monotone formulas depending on literals $x, \bar{x}$; in other words, for each variable x , we replace $\neg x$ by $\bar{x}$. The double sequent

$$\text{CNF}(\varphi) \longleftrightarrow C(\varphi; \delta)$$

has a tree-like proof of size $2^{O(2^D)}$ cutting only on literals.

Proof. Since φ has at most 2^D leaves, it depends on at most 2^D variables, and twice as many literals. Lemma 2 provides the necessary proof of size $2^{O(D)} 2^{2^{D+1}} = 2^{O(2^D)}$.

Lemma 11. Let φ be a formula, and $\mathbf{d} = d_1, \dots, d_k$ be a vector of increasing positive integers. The double sequent

$$C(\varphi; \mathbf{d}) \longleftrightarrow D(\varphi; \mathbf{d})$$

has a tree-like proof of size $2^{O(k2^x)}$ cutting only on formulas of depth at most $k + 2$, where $x = \text{ex}(\varphi; \mathbf{d})$.

Proof. The proof is by induction on k . When $k = 0$, we prove that the CNF and DNF forms are equivalent using Lemma 2.

If $k > 0$, let ψ be the part of the φ up to depth d_1 . According to the recipe of Definition 4, $C(\varphi; \mathbf{d})$ is obtained from $\text{CNF}(\psi)$ by substituting equivalent forms v_C^+, v_C^- of subformula literals $v, \bar{v}$. Similarly, $D(\varphi; \mathbf{d})$ is obtained from $\text{DNF}(\psi)$ by substituting equivalent forms v_D^+, v_D^- of subformula literals $v, \bar{v}$.

Note that v_C^+, v_D^+ are disjunctive and conjunctive flattened forms of the same subformula of φ . Similarly, v_C^-, v_D^- are obtained by dualizing conjunctive and disjunctive normal forms of the same subformula of φ . Therefore, the induction hypothesis gives us proofs of the double sequents $v_C^+ \longleftrightarrow v_D^+$ and $v_C^- \longleftrightarrow v_D^-$ (the latter, using Lemma 6). Repeated applications of Lemma 4, together with cuts on the ‘hybrids’, allow us to lift at most $2^{O(k2^x)}$ instances of these proofs to a proof of $C(\varphi; \mathbf{d}) \longleftrightarrow \text{CNF}(\psi)(v_D^\pm)$; see the proof of Lemma 9 for more details.

Substituting $v_D^\pm$ in the proof of $\text{CNF}(\psi) \longleftrightarrow \text{DNF}(\psi)$ using Lemma 3, we obtain a proof of the double sequent $\text{CNF}(\psi)(v_D^\pm) \longleftrightarrow D(\varphi; \mathbf{d})$. The proof is completed by an application of the cut rule.