Cut Elimination *In Situ*

Revised version. Comments appreciated.

Sam Buss*
Department of Mathematics
University of California, San Diego
La Jolla, CA 92093-0112, USA
sbuss@math.ucsd.edu

December 5, 2013

Abstract

We present methods for removing top-level cuts from a sequent calculus or Tait-style proof without significantly increasing the space used for storing the proof. For propositional logic, this requires converting a proof from tree-like to dag-like form, but it most doubles the number of lines in the proof. For first-order logic, the proof size can grow exponentially, but the proof has a succinct description and is polynomial-time uniform. We use direct, global constructions that give polynomial time methods for removing all top-level cuts from proofs. By exploiting prenex representations, this extends to removing all cuts, with final proof size near-optimally bounded superexponentially in the alternation of quantifiers in cut formulas.

1 Introduction

Gentzen's technique of cut elimination, together with the closely related normalization, is arguably the most important construction of proof theory. The importance of cut elimination lies partly in its connections to constructivity, and indeed cut elimination is algorithmic and can be carried out effectively. The present paper focuses on algorithms for cut elimination in the setting of pure propositional logic and pure first-order logic. We introduce methods for removing top-level cuts from a proof without significantly increasing the

*Supported in part by NSF grants DMS-1101228 and CCF-1213151 and by a grant from the Simons Foundation (#208717 to Sam Buss).

space used for generating the proof. Of course, it is well-known that eliminating top-level cuts can make proof size grow exponentially, so it requires some special care to describe the resulting proof without any significant increase in space. For propositional logic, our methods require converting a proof from tree-like to dag-like form, but at most double the number of lines in the proof. For first-order logic, the proof size can grow exponentially; in fact, both the number of lines in the proof and the size of the terms can grow exponentially. However, our constructions give polynomial size dag representations for the terms, and succinct descriptions of the proof that give a polynomial time uniform description of the proof and its terms.

Along with the small space usage, our cut elimination methods give direct, global constructions. We define direct, concrete descriptions of the proof that results from eliminating the top-level cuts. Our construction is "global" in that it operates on the entire proof and eliminates all top-level cuts at once.

Our constructions synthesize and generalize a number of prior results from proof complexity and continuous cut elimination. Our immediate motivation arose from the desire to find global versions of the polynomial time algorithms for the continuous cut elimination used by Aehlig-Beckmann [1] and Beckmann-Buss [4]. Continuous cut elimination was developed by Mints [12, 11] for the analysis of higher order logics, and [1] introduced its use for the analysis of bounded arithmetic. In particular, [1, 4] required polynomial time constructions of proofs. Like Mints, they create proofs step-by-step and use a special Rep (for "repetition" or "repeat") inference to slowdown the construction of proofs. In contrast, we shall give direct (not step-by-step) constructions, and avoid the use of a Rep inference.

There is extensive prior work giving upper bounds on the complexity of cut elimination in propositional and first-order logic, including [13, 14, 8, 6, 17, 18, 7, 16, 5, 2]. Some of the best such bounds measure the complexity of proofs in terms of the height of proofs [13, 17, 18, 7, 16, 5]. Loosely speaking, these results work by removing top-level connectives from cut formulas, at the cost of exponentiating the height of the proof, and repeating this to remove all cuts from a proof. Zhang [17] and Gerhardy [7] bound the height of cut free proofs in terms of the nesting of quantifiers in cut formulas; namely, if quantifiers are nested to depth d without any intervening propositional connectives, then cut elimination requires a height increase of only an exponential stack of 2's of height $d+2$. They further show that cut-elimination can remove a top-level block of $\exists$ and $\vee$ (respectively, $\forall$ and $\wedge$) connectives at the cost of a single exponential increase in proof height.

In contrast, the present paper works with proof size rather than proof height. Somewhat counterintuitively, blocks of arbitrarily nested $\exists$ and $\wedge$ connectives (respectively, $\forall$ and $\vee$ connectives) can be removed all at once, with a single exponential increase in proof size.

Krajíček [9, 10], Razborov [15], and Beckmann-Buss [3] have given complexity bounds for reducing the depth (alternation of $\vee$'s and $\wedge$'s) of formulas in constant depth propositional Frege or Tait-style proofs. Reducing the depth of formulas in a proof is essentially equivalent to removing the outermost blocks of like (propositional) connectives from cut formulas. Krajíček [9] and later Beckmann-Buss [3] show that the depth of formulas in a constant depth proof can be reduced from $d+1$ to d at the cost of converting the proof from tree-like format to dag-like format with only a polynomial increase in proof size. Our Theorem 3 below is similar to Lemma 6 of [3] in this regard, but gives a more explicitly uniform construction, and works even if there are multiple nested outermost like quantifiers that need to be eliminated.

This paper deals with cut elimination for a Tait-style calculus instead a Gentzen sequent calculus. In the setting of classical logic, our results all apply immediately to cut elimination in a Gentzen sequent calculus.[1] We assume the reader has some familiarity with sequent calculi or Tait calculi, but Section 2 begins with formal definitions of our Tait-style proof system, including definitions of proof size and cut formula complexity. It also describes the basic ideas behind the later constructions. Section 3 shows that, for tree-like propositional proofs, outermost like connectives in cut formulas can be removed at the cost of converting the proof to dag-like form, while at most doubling the number of lines in the proof. Sections 4 and 5 extend this to first-order logic, but now, instead of forming a dag-like proof of the same size, the number of lines in the proof can become exponentially larger. However, the exponentially long proof still has a direct, global, polynomial-time specification. For expository purposes, Section 4 first shows how to eliminate all top-level like quantifiers from cut formulas. Section 5 then combines the earlier constructions to show how to eliminate all outermost $\forall$ and $\vee$ connectives. In light of the duality of the Tait calculus, this is the same as removing all top-level $\exists$ and $\wedge$ connectives. Our constructions use direct methods that reduce the cut-formula complexity for multiple cuts simultaneously.

So far, we have discussed only the problem of removing the top-level

[1]Tait systems do not work as well as the Gentzen sequent calculus for non-classical systems such as intuitionistic logic. Thus our results would need to be modified to apply to intuitionistic logic, for instance.

connectives from cut formulas. Obviously, the process could be iterated to remove all cuts. Define the *alternating quantifier depth* of a formula as the maximum number of alternating blocks of existential and universal quantifiers along any branch in the tree representation of the formula (with negations pushed to the atoms, but allowing $\land$ and $\lor$ connectives to appear arbitrarily along the branch). Let $\mathsf{aqd}(P)$ be the maximum alternating quantifier depth of any cut formula in the proof P. Section 6 proves that it is possible to convert P into a cut free proof of the same end cedent, with the size of P bounded by $2_d^{|P|}$ for $d = \mathsf{aqd}(P) + O(1)$. Here $|P|$ is the number of lines in P, and the superexponential function 2_d^a is defined by $2_0^a = a$ and $2_{i+1}^a = 2^{2_i^a}$. This improves on what can be obtained straightforwardly using the constructions of Sections 3-5 or from the prior bounds obtained by Zhang [17], Gerhardy [7], and Beckmann-Buss [5], since we bound the height of the stack of two's in terms of the number of alternations of quantifiers without regard to intervening $\land$'s or $\lor$'s. The basic idea for the proof in Section 6 is to first modify P so that all cut formulas are in prenex form, and then apply the results of Section 4. The results of Section 6 do not depend on either Section 3 or 5; but we do appeal to constructions of [17, 7, 5] to handle removing cuts on quantifier free formulas.

2 Preliminaries

2.1 Tait calculus

Our first-order Tait system uses logical connectives $\land$, $\lor$, $\exists$ and $\forall$, and a language of function symbols, constant symbols, and predicate symbols. Terms and atomic formulas are defined as usual. A *literal* is either an atomic formula $P(\vec{s})$ or a negated atomic formula $\overline{P(\vec{s})}$. Formulas are formed using connectives $\land$, $\lor$, $\forall$ and $\exists$. The negation of complex formulas is inductively defined by defining $\overline{(\overline{p})}$, $\overline{B \land C}$, $\overline{B \lor C}$, $\overline{(\exists x)A}$, and $\overline{(\forall x)A}$ to be the formulas p, $\overline{B} \lor \overline{C}$, $\overline{B} \land \overline{C}$, $(\forall x)\overline{A}$, and $(\exists x)\overline{A}$, respectively.

We adopt a convention from the Gentzen sequent calculus and assume that first-order variables come in two sorts: free variables (denoted with letters $a, b, c, \ldots$) and bound variables (denoted with letters $x, y, \ldots$). Free variables cannot be quantified and must appear only freely. A bound variable x may occur in formulas only within the scope of a quantifier $(\forall x)$ or $(\exists x)$ that binds it.

A line of a Tait calculus proof, called a *cedent*, consists of a set of formulas. The intended meaning of a cedent is the disjunction of its members. The allowable rules of inference are shown in Figure 1. It should be noted that an initial cedent $A, \overline{A}$ must have A atomic. We allow Tait proofs to be either tree-like or dag-like. The usual conditions for eigenvariables apply to $\forall$ inferences. The formulas introduced in the lower cedents of inferences are called the *principal* formula of the inference: these are the formulas $A \wedge B$, $A \vee B$, $(\exists x)A(x)$ and $(\forall x)A(x)$ in Figure 1. The formulas eliminated from the upper cedent are called *auxiliary* formulas: these are the formulas A, B, $A(s)$, $A(b)$, A, and $\overline{A}$ in the figure. The auxiliary formulas of a cut inference are called *cut formulas*. Formulas that appear in the sets Γ and Γ_i are called *side* formulas.

The $\wedge$ and cut inferences have two cedents as hypotheses, which are designated the *left* and *right* upper cedents. For a cut inference, we require that the outermost connective of the left cut formula A not be an $\wedge$ or $\exists$ connective; equivalently, the outermost connective of the right cut formula $\overline{A}$ is not $\vee$ or $\forall$. This restriction on A's outermost connective causes no loss of generality, since the order of the upper cedents can always be reversed. (We sometimes display cuts with upper cedents out of order, however.) For an $\wedge$ inference, the left-right order of the upper cedents is dictated by the order of the conjunction; except in the case where A and B are the same formula, and then the upper cedents are put in some arbitrary left-right order.

The left-right ordering of upper cedents allows us to define the post-ordering of the cedents of a tree-like proofs. The *postordering* of the nodes of a tree T is the order of the nodes output by the following recursive traversal algorithm: Starting at the root of T, the traversal algorithm first recursively traverses the child nodes in left-to-right order, and then outputs the root node. The postorder traversal of the underlying proof tree induces an ordering of the cedents in the proof.

Axioms (initial cedents) and weakening inferences are ignored when measuring the size or height of P. Thus, the *size*, $|P|$, of a Tait proof P is defined as the number of $\vee$, $\wedge$, $\forall$, $\exists$, and cut inferences in P. The *height*, $h(P)$, of P is the maximum number of these kinds of inferences along any branch of P.

The fact that cedents are *sets* rather than multisets or sequences means that if a formula is written twice on a line, it appears only once in the cedent. For instance, in the $\vee$ inference, is it possible that $A \vee B$ is a member of Γ. It is also possible that A (say) appears in Γ, in which case both A and $A \vee B$ appear in the conclusion of the inference. This latter possibility, however, makes our analysis of cut elimination more awkward, since we will track occurrences of formulas along paths in the proof tree. The problem is that there will be an ambiguity about how to track the formula A in the case where it "splits into two", for example in an $\vee$ inference by both being a member of Γ and being used to introduce $A \vee B$. The ambiguity

$$\text{Axiom: } A, \overline{A} \qquad \text{Weakening: } \frac{\Gamma}{\Gamma, \Delta}$$

$$\wedge\text{: } \frac{A, \Gamma_1 \qquad B, \Gamma_2}{A \wedge B, \Gamma_1, \Gamma_2} \qquad \vee\text{: } \frac{A, B, \Gamma}{A \vee B, \Gamma}$$

$$\exists\text{: } \frac{A(s), \Gamma}{(\exists x)A(x), \Gamma} \qquad \forall\text{: } \frac{A(b), \Gamma}{(\forall x)A(x), \Gamma}$$

$$\text{Cut: } \frac{A, \Gamma_1 \qquad \overline{A}, \Gamma_2}{\Gamma_1, \Gamma_2}$$

Figure 1: The rules of inference for a Tait system. The lines of the proof are to be interpreted as sets of formulas. The formula A of the *axiom* rule must be atomic. The free variable b of the $\forall$ inference is called an *eigenvariable* and may not occur in the lower cedent.

can be avoided by considering proofs that satisfy the following "auxiliary condition":

Definition A Tait proof P satisfies the *auxiliary condition* provided that no inference has an auxiliary formula also appearing as a side formula. Specifically, referring to Figure 1, the auxiliary condition requires the following to hold:

a. In an $\vee$ inference, neither A nor B may occur in Γ.

b. In an $\wedge$ inference, neither A nor B may occur in Γ_1 or Γ_2.

c. In an $\exists$ inference, $A(s)$ may not occur in Γ.

d. In a cut inference, neither A nor $\overline{A}$ may occur in Γ_1 or Γ_2.

Note that the eigenvariable condition already prevents $A(b)$ from occurring in the side formulas of a $\forall$ inference.

Lemma 1 *Let P be a [tree-like] Tait proof. Then there is a [tree-like] Tait proof P' satisfying the auxiliary condition proving the same conclusion as P. Furthermore, $|P'| \leq |P|$ and $h(P') \leq h(P)$.*

The proof of the lemma is straightforward using the fact that weakening inferences do not count towards proof size or height.

A *path* in a proof P is a sequence of one or more cedents occurring in P, with the $(i+1)^{st}$ cedent a hypothesis of the inference inferring the i^{th} cedent, for all i. A *branch* is a path that starts at the conclusion of P and ends at an initial cedent.

Suppose P is tree-like and satisfies the auxiliary condition. Also suppose a formula A occurs in two cedents C_1 and C_2 in P, and let A_1 and A_2 denote the *occurrences* of A in C_1 and C_2, respectively. We call A_1 a *direct ancestor* of A_2 (equivalently, A_2 is a *direct descendant* of A_1) provided there is a path in P from C_2 to C_1 such that the formula A appears in every cedent in the path.[2] If A_1 is the principal formula of an inference, or occurs in an axiom, that we say A_1 is a place where A_2 is *introduced*. If A_2 is an auxiliary formula, then we say A_2 is the place where A_1 is *eliminated*. In view of the tree-like property of P, every formula occurring in P either has a unique place where it is eliminated or has a direct descendant in the conclusion of P. However, due to the implicit use of contraction in the inference rules, formulas occurring in P may be introduced in multiple places.

The notions of direct descendant and direct ancestor can be generalized to "descendant" and "ancestor" by tracking the flow of subformulas in a proof. If $\mathcal{I}$ is an $\land$, $\lor$, $\exists$, or $\forall$ inference, then the principal formula of $\mathcal{I}$ is the (only) *immediate descendant* of each auxiliary formula of $\mathcal{I}$. Then, the "descendant" relation is the reflective, transitive closure of the union of the immediate descendant and direct descendant relations. Namely, a formula A' occurring in P is a *descendant* of a formula A occurring in P iff there is a sequence of formula occurrences in P, starting with A and ending with A' such that each formula in the sequence is the immediate descendant or a direct descendant of the previous formula in the sequence. We also call A an *ancestor* of A'.

The definitions of descendant and ancestor apply to formulas that appear in cedents. Similar notions also apply to subformulas. Suppose A and B are formulas appearing in cedents with B a descendant of A. Let C be a subformula of A. We wish to define a unique subformula D of B, such that C *corresponds to* D. This unique subformula is intended to be defined in the obvious way, with each subformula in an upper cedent of an inference corresponding to a subformula in the lower sequent. Assume P is tree-like and satisfies the auxiliary condition. The "corresponds" relation is defined by taking the reflexive, transitive closure of the following conditions.

- The formula $A(s)$ in an $\exists$ inference corresponds to the subformula $A(x)$ in the lower sequent.

- In a $\forall$ inference, the formula $A(b)$ corresponds to the subformula $A(x)$.

[2] This definition works because P satisfies the auxiliary condition.

- In an $\land$ or $\lor$ inference, the formulas A and B in the upper cedent(s) correspond to the subformulas A and B shown in the lower cedent. Except for an $\lor$ inference in which A and B are the same formula, the auxiliary formula corresponds to the subformula denoted A in the lower cedent. That is, in this case, the $\lor$ inference is treated as if it were defined as

$$\frac{A, \Gamma}{A \lor B, \Gamma}$$

- If C is a subformula of a side formula, namely of a formula A in Γ, Γ_1, or Γ_2 in Figure 1, then C corresponds to the same subformula of the occurrence of A in the lower cedent.

- If A and B appear in the upper and lower cedent of an inference and A corresponds to B and if C is the i-th subformula of A, then C corresponds to the i-th subformula D of B, where the subformulas of C and D are ordered (say) according to the left-to-right positions of their principal connectives.

It is often convenient to assume proofs use free variables in a controlled fashion. The following definition is slightly weaker than the usual definition, but suffices for our purposes.

Definition A proof P is in *free variable normal form* provided that each free variable b is used at most once as an eigenvariable, and provided that when b is used as an eigenvariable for inference $\mathcal{I}$, then b appears in P only above $\mathcal{I}$ (that is, each occurrence of b occurs in a cedent reachable from $\mathcal{I}$ by some path in P). The variables c that appear in P but are not used as eigenvariables are called the *parameter variables* of P.

Any tree-like proof P may be put into free variable normal form without increasing its size or height; furthermore, this can be done while enforcing the auxiliary condition.

2.2 The basic constructions

This section describes the basic ideas and constructions used for the cut-elimination results obtained in Sections 3 and 4.

The first important tool is a generalization of the well-known inversion lemmas for the outermost $\forall$ and $\lor$ connectives of a formula. Assume we have a tree-like proof P, in free variable normal form, that ends with the cedent $\Gamma, A \vee B$. Then there is a proof P' of Γ, A, B, with P' also tree-like, and with $|P'| \leq |P|$ and $h(P') \leq h(P)$. Similarly, if P ends with $\Gamma, (\forall x)A(x)$ and t is any term, then there is a proof P'' of $\Gamma, A(t)$, with P'' also tree-like and satisfying the same conditions on its size and height. The proofs are quite simple: P' is obtained from P by replacing all direct ancestors of $A \vee B$ with A, B and removing all $\vee$ inferences that introduce a direct ancestor of $A \vee B$. Likewise, if t does not contain any eigenvariables of P, then P'' is formed by replacing all direct ancestors of $(\forall x)A(x)$ with $A(t)$, and removing the $\forall$ inferences that introduce these direct ancestors and replacing their eigenvariables with t.

Iterating this construction allows us to formulate an inversion lemma that works for the entire set of outermost $\vee$ and $\forall$ connectives. If B is a subformula of A, we call B an $\vee\forall$-*subformula* of A if every connective of A containing B in its scope is an $\vee$ or a $\forall$. Similarly, a connective $\vee$ or $\forall$ is said to be $\vee\forall$-*outermost* if it is not in the scope of any $\exists$ or $\wedge$ connective. Let P be a tree-like proof of Γ, A, and let $B_1, \ldots, B_k$ enumerate the minimal $\vee\forall$-subformulas of A in left-to-right order. The subformulas B_i are called the $\vee\forall$-*components* of A. Note that each B_i is atomic or has as outermost connective an $\wedge$ or an $\exists$.

Lemma 2 *Let P, A, $B_1, \ldots, B_k$ be as above. Let σ be any substitution mapping free variables to terms. Then there is a proof P' of $\Gamma\sigma, B_1\sigma, \ldots, B_k\sigma$ such that P' is tree-like and $|P'| \leq |P|$ and $h(P') \leq h(P)$.*

The lemma is proved by iterating the inversion lemmas for $\vee$ and $\forall$.

Section 3 will give the details how to simplify cuts in a propositional Tait calculus proof by removing all outermost $\vee$ (or, all outermost $\wedge$) connectives from cut formulas. As a preview, we give the idea of the proof, which depends on the inversion lemma for $\vee$. Namely, suppose the proof P ends with a cut on the formula $A \vee (B \vee C)$, as shown in Figure 2. The right cut formula, in the final line of the subproof R, is in the dual form $\overline{A} \wedge (\overline{B} \wedge \overline{C})$ of course. Now suppose that in the subproof R there are the two pictured $\wedge$ inferences that introduce the formulas $(\overline{B} \wedge \overline{C})$ and then $\overline{A} \wedge (\overline{B} \wedge \overline{C})$.

By the inversion lemma for $\vee$, the proof Q can be transformed into a proof Q' of A, B, C, Γ_1 with no increase in size or height. The cut in P can thus be removed by replacing the $\wedge$ inferences in R with cuts to obtain the proof P' shown in Figure 3. Note that this has replaced the $\wedge$ inference introducing $\overline{B} \wedge \overline{C}$ with two cuts, one on B and one on C, and replaced the $\wedge$ inference introducing $\overline{A} \wedge (\overline{B} \wedge \overline{C})$ with a cut on A. Overall, two

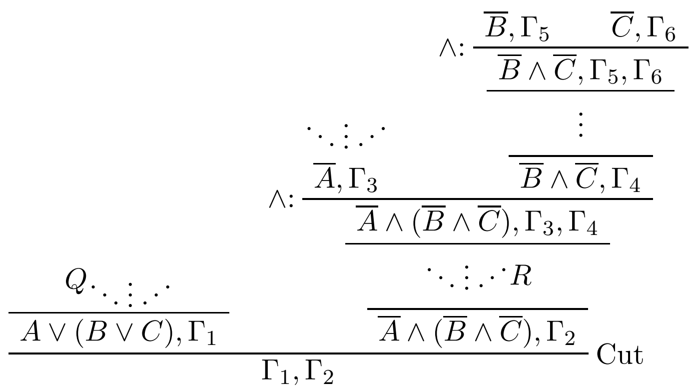

Figure 2: A simple example of $\vee$ cut to be eliminated. Q and R are the subproofs deriving the hypotheses of the cut.

$\wedge$ inferences and one cut inference in P have been replaced by three cut inferences in P'. More generally, due to contractions, there can be $k_1 \geq 1$ inferences in P that introduce $\overline{B} \wedge \overline{C}$, and $k_2 \geq 1$ inferences that introduce $\overline{A} \wedge (\overline{B} \wedge \overline{C})$: these $k_1 + k_2$ many $\wedge$ inferences and the cut inference in P are replaced by $2k_1 + k_2$ many cut inferences in P'. Thus the size of P' is no more than twice the size of P. The catch though, is that P' may now be dag-like rather than tree-like.

Finally, it should be noted that P' is obtained from P by moving the subproof Q' and the subproof deriving $\overline{A}, \Gamma_3$ "rightward and upward" in the proof. This is crucial in allowing us to remove multiple cuts at once. Intuitively, the final cut of P plus all the cuts that lie in the subproofs Q or R can be simplified in parallel without any unwanted "interference" between the different cuts.

Figure 4 shows a proof P from which the outermost $\vee$ and $\forall$ (dually, $\wedge$ and $\exists$) connectives can be removed from cut formulas. The left subproof Q can be inverted to give a proof Q' of $A(r,s), B(r,t), \Gamma_1$, and this is used to form the proof P' shown in Figure 5. In this simple example, an $\wedge$ inference, three $\exists$ inferences, and the cut inference are replaced by just two cut inferences. As in the $\vee$ example, the proof P' is formed by moving (instantiations of) subproofs of P rightward. In particular, the subproof in P ending with $(\exists y)\overline{A(r,y)}, \Gamma_3$ has become a proof of $B(r,t), \Gamma_1, \Gamma_3$ and has been moved rightward in the proof so as to be cut against $\overline{B(r,t)}, \Gamma_6$.

The general case of removing quantifiers is more complicated however. For instance, there might be multiple places where the formula $(\exists y)\overline{A(x,y)}$

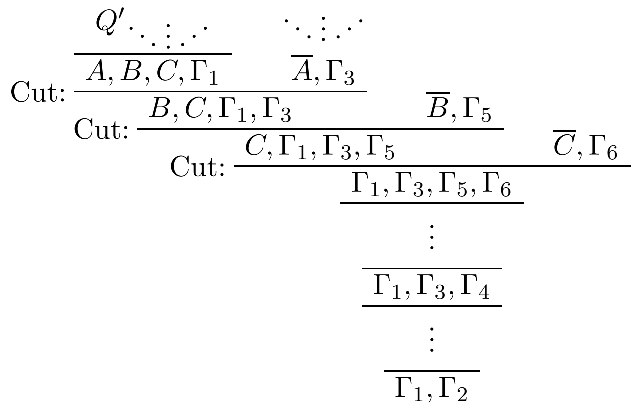

Figure 3: The proof P' obtained after eliminating the cut of Figure 2.

is introduced, using k_1 terms $s_1, \ldots, s_{k_1}$. Likewise, there could be k_2 terms t_j used for introducing the formula $(\exists y)\overline{B(x,y)}$, and k_3 terms r_ℓ for introducing the $(\exists x)$. In this case we would need $k_1 k_2 k_3$ many inversions of Q, namely, proofs $Q_{i,j,\ell}$ of $A(r_\ell, s_i), B(r_\ell, t_j), \Gamma_1$ for all $i \le j_1$, $j \le k_2$, and $\ell \le k_3$. The result is that P' can have size exponential in the size of P; there is, however, still a succinct description of P' which can be obtained directly from P. This will be described in Section 4.

3 Eliminating like propositional connectives

This section describes how to eliminate an outermost block of propositional connectives from cut formulas. The construction applies to proofs in first-order logic.

Definition Suppose B is a subformula occurring in A. Then B is an $\vee$-*subformula* of A iff B occurs in the scope of only $\vee$ connectives. The notion of $\wedge$-*subformula* is defined similarly.

An $\vee$-*component* (resp., $\wedge$-*component*) of A is a minimal $\vee$-subformula (resp., $\wedge$-subformula) of A.

Definition An $\vee/\wedge$-*component of a cut formula in* P is an $\vee$-component of a left cut formula in P or an $\wedge$-component of a right cut formula in P.

Theorem 3 *Let P be a tree-like Tait calulus proof of Γ. Then there is a dag-like proof P', also of Γ, such that each cut formula of P' is an $\vee/\wedge$-component*

$$
\dfrac{
\dfrac{\stackrel{Q}{\vdots}}{(\forall x)((\forall y)A(x,y) \lor (\forall y)B(x,y)), \Gamma_1}
\qquad
\dfrac{
\stackrel{\vdots}{
\exists\!:\ \dfrac{
\wedge\!:\ \dfrac{
\dfrac{\stackrel{\vdots}{\exists\!:\ \dfrac{\overline{A(r,s)}, \Gamma_5}{(\exists y)\overline{A(r,y)}, \Gamma_5}}}{(\exists y)\overline{A(r,y)}, \Gamma_3}
\qquad
\dfrac{\stackrel{\vdots}{\exists\!:\ \dfrac{\overline{B(r,t)}, \Gamma_6}{(\exists y)\overline{B(r,y)}, \Gamma_6}}}{(\exists y)\overline{B(r,y)}, \Gamma_4}
}{(\exists y)\overline{A(r,y)} \land (\exists y)\overline{B(r,y)}, \Gamma_3, \Gamma_4}
}{(\exists x)((\exists y)\overline{A(x,y)} \land (\exists y)\overline{B(x,y)}), \Gamma_3, \Gamma_4}
}
}{(\exists x)((\exists y)\overline{A(x,y)} \land (\exists y)\overline{B(x,y)}), \Gamma_2}
}{\Gamma_1, \Gamma_2}\ \text{Cut}
$$

Figure 4: A simple example of cuts using $\lor$ and $\exists$ to be eliminated.

of a cut formula of P, and such that $|P'| \leq 2 \cdot |P|$ and hence $h(P') \leq 2 \cdot |P|$. Furthermore, given P as input, the proof P' can be constructed by a polynomial time algorithm.

Note that P' is obtained by simplifying *all* the cut formulas in P that have outermost connective $\land$ or $\lor$.

Without loss of generality, by Lemma 1, P satisfies the auxiliary condition. The construction of P' depends on classifying the formulas appearing in P according to how they descend to cut formulas. For this, each formula B in P can be put into exactly one of the following categories (α)-(γ).

(α) B has a left cut formula A as a descendant and corresponds to an $\lor$-subformula of A, or

(β) B has a right cut formula A as a descendant and corresponds to an $\land$-subformula of A, or

(γ) Neither (α) nor (β) holds.

Definition Let B be an occurrence of a formula in P, and suppose B is in category (β) with a cut formula A as a descendant. The formula A is a conjunction $\bigwedge_{i=1}^{k} C_i$ of its $k \geq 1$ many $\land$-components (parentheses are suppressed in the notation). The formula B is a subconjunction of A of the form $\bigwedge_{i=m}^{\ell} C_i$ where $1 \leq m \leq \ell \leq k$. The $\land$-components of A to the right of B are $C_{\ell+1}, \ldots, C_k$. The negations of these, namely $\overline{C}_{\ell+1}, \ldots, \overline{C}_k$, are called the *pending implicants for B*.

$$
\text{Cut: }\dfrac{\dfrac{Q'\;\ddots\vdots\,\cdot^{\cdot^{\cdot}}}{A(r,s),B(r,t),\Gamma_1} \qquad \overline{A(r,s)},\Gamma_5}{\dfrac{B(r,t),\Gamma_1,\Gamma_5}{\vdots}}
$$

$$
\text{Cut: }\dfrac{\overline{B(r,t),\Gamma_1,\Gamma_3} \qquad \overline{B(r,t)},\Gamma_6}{\dfrac{\Gamma_1,\Gamma_3,\Gamma_6}{\dfrac{\vdots}{\dfrac{\Gamma_1,\Gamma_3,\Gamma_4}{\dfrac{\vdots}{\Gamma_1,\Gamma_2}}}}}
$$

Figure 5: The results of eliminating the cuts in Figure 4.

Each formula B in P will be replaced by a cedent denoted $*(B)$. For B in category (α), $*(B)$ is the cedent consisting of the $\lor$-components of B. For B in category (β), $*(B)$ is the (possibly empty) cedent containing the pending implicants for B. For B in category (γ), $*(B)$ is the cedent containing just the formula B.

Definition The *jump target* of a category (β) formula B in P is the first cut or $\land$ inference below the occurrence of B which has some descendant of B as an auxiliary formula in its right upper cedent. The jump target will be either:

$$
\dfrac{\overline{D},\Gamma_1 \qquad D,\Gamma_2}{\Gamma_1,\Gamma_2} \qquad \text{or} \qquad \dfrac{C,\Gamma_1 \qquad D,\Gamma_2}{C\land D,\Gamma_1,\Gamma_2} \tag{1}
$$

where the formula D is either equal to B (a direct descendant of B) or is of the form $((\cdots(B\land B_1)\land\cdots\land B_{k-1})\land B_k)$ with $k\geq 1$ (since only $\land$ inferences can operate on B until reaching the jump target). The left upper cedent of the jump target (that is, $\overline{D},\Gamma_1$ or C,Γ_1) is called the *jump target cedent.* The auxiliary formula of the left upper cedent, that is $\overline{D}$ or C, is called the *jump target formula.*

We shall consistently suppress parentheses when forming disjunctions and conjunctions. For instance, the formula $((\cdots(B\land B_1)\land\cdots\land B_{k-1})\land B_k)$ above would typically be written as just $B\land B_1\land\cdots\land B_k$. It should be clear from the context what the possible parenthesizations are.

Lemma 4 *Suppose B is category (β) formula in P. Let $C_1, \ldots, C_k$ be as above, so $B = \bigwedge_{i=m}^{\ell} C_i$ and the pending implicants of B are $\overline{C}_{\ell+1}, \ldots, \overline{C}_k$. Consider B's jump target, namely one of the inferences shown in (1), and let E be the jump target formula, that is, either $\overline{D}$ or C. Then $*(E)$ is equal to the cedent $\overline{C}_m, \ldots, \overline{C}_k$.*

Proof If the jump target of B is a cut inference, then D is $\bigwedge_{i=1}^{k} C_i$. In this case, $E = \overline{D}$ is the formula $\overline{C}_1 \vee \cdots \vee \overline{C}_k$, and $m = 1$. It follows that E is category (α), and $*(E) = \overline{C}_1, \ldots, \overline{C}_k$, so the lemma holds. On the other hand, if the jump target is an $\wedge$ inference, then D equals $C_m \wedge \cdots \wedge C_r$ for some $r \leq k$, and $E = C$ equals $\bigwedge_{i=j}^{m-1} C_i$ for some $1 \leq j < m$. In this case, E is category (β), and $*(E)$ again equals $\overline{C}_m, \ldots, \overline{C}_k$. □

Proof (of Theorem 3). The cedents of P' are formed by modifying each cedent Δ of P to form a new cedent Δ^*, called the $*$-*translation* of Δ. A formula B *occurs in or below* Δ if it is in Δ or is in some cedent below Δ in P. For each Δ in P, the cedent Δ^* is defined to include the formulas $*(B)$ for all formulas B which occur in or below Δ.

Theorem 3 is proved by showing that the cedents Δ^* can be put together to form a valid proof P'. This requires making the following modifications to P: (1) For any inference in P that introduces an $\wedge$-component of a right cut formula, we must insert at that point in P' a cut on that $\wedge$-component using (the $*$-translation of) its jump target cedent. (2) When forming P', we remove from P every $\wedge$ inference that introduces an $\wedge$-subformula of a right cut formula, every $\vee$ inference that introduces an $\vee$-subformula of a left cut formula, and every cut inference of P. (3) Weakening inferences are added as needed. These changes are described in detail below, where we describe how to combine the cedents Δ^* to form the proof P'. We consider separately each possible kind of inference in P.

For the first case, consider the case where Δ is an initial cedent $B, \overline{B}$. (Surprisingly, this is the hardest case of the proof.) Our goal is to show how the cedent Δ^* is derived in P'. As a first subcase, suppose neither B nor $\overline{B}$ is in category (β), so neither descends to an $\wedge$-component of a right cut formula. Since B is atomic, and B and $\overline{B}$ are each in category (α) or (γ), we have $*(B) = B$ and $*(\overline{B}) = \overline{B}$, respectively. The cedent Δ^* is equal to $B, \overline{B}, \Lambda$, where Λ contains the formulas $*(E)$ for all formulas E that occur below the cedent $B, \overline{B}$. The proof P' merely derives $B, \overline{B}, \Lambda$ from $B, \overline{B}$ by a weakening inference. (Recall that weakening inferences do not count towards the size or height of proofs.)

For the second subcase, suppose exactly one of B and $\overline{B}$ are in category (β). Without loss of generality, we may assume B is of category (β), and $\overline{B}$ is not. The formula B descends to a right cut formula $\bigwedge_{i=1}^{k} B_i$, and corresponds uniquely to one of its $\wedge$-components B_ℓ. We have $1 \leq \ell \leq k$, and $\overline{B}_{\ell+1}, \ldots, \overline{B}_k$ are the pending implicants of $B = B_\ell$. Since $\overline{B}$ is atomic and not in category (β), $*(\overline{B}) = \overline{B} = \overline{B}_\ell$. Thus, Δ^* is equal to

$$\overline{B}_{\ell+1}, \ldots, \overline{B}_k, \overline{B}_\ell, \Lambda. \tag{2}$$

As before, the cedent Λ is the set of *-translations of formulas that appear below Δ in P.

The jump target for B has the form

$$\text{Cut:}\ \frac{\overline{D}, \Gamma_1 \qquad D, \Gamma_2}{\Gamma_1, \Gamma_2} \qquad \text{or} \qquad \wedge\text{:}\ \frac{C, \Gamma_1 \qquad D, \Gamma_2}{C \wedge D, \Gamma_1, \Gamma_2} \tag{3}$$

By Lemma 4, the $*$-translation of the upper left cedent has the form

$$\overline{B}_\ell, \ldots, \overline{B}_k, \Lambda' \tag{4}$$

where Λ' contains the formulas $*(E)$ for all formulas E occurring in or below the lower cedent of the inference (3). Of course, $\Lambda' \subseteq \Lambda$. Thus, in P', the cedent Δ^* is derived from the cedent (4) by a weakening inference.

In the third subcase, both B and $\overline{B}$ are in category (β). As in the previous subcase, $*(B)$ has the form $\overline{B}_{\ell+1}, \ldots, \overline{B}_k$, and the $*$-translation of its jump target cedent has the form

$$\overline{B}_\ell, \ldots, \overline{B}_k, \Lambda'$$

with $B_\ell = B$. Likewise, $*(\overline{B})$ has the form $B'_{\ell'+1}, \ldots, B'_{k'}$ and the $*$-translation of $\overline{B}$'s jump target cedent has the form

$$\overline{B}'_{\ell'}, \ldots, \overline{B}'_{k'}, \Lambda''$$

where $B'_{\ell'} = \overline{B}$. These two cedents combine with a cut on the formula B to yield the inference

$$\frac{\overline{B}_\ell, \ldots, \overline{B}_k, \Lambda' \qquad \overline{B}'_{\ell'}, \ldots, B'_{k'}, \Lambda''}{\overline{B}_{\ell+1}, \ldots, \overline{B}_k, \overline{B}'_{\ell'+1}, \ldots, \overline{B}'_{k'}, \Lambda', \Lambda''}$$

Since $\Lambda', \Lambda'' \subseteq \Lambda$, the cedent Δ^* is derivable with one additional weakening inference. This completes the argument for the case of an initial cedent.

Note that in the first two subcases, the initial cedent is eliminated, while bypassing a cut or $\wedge$ inference. In the third subcase, the initial cedent is replaced with a cut inference on an atomic formula.

For the second case of the proof of Theorem 3, consider a weakening inference

$$\frac{\Gamma}{\Gamma, \Delta}$$

in P. Here, the upper and lower sequents have exactly the same $*$-translations; that is, Γ^* is the same as $(\Gamma, \Delta)^*$. Thus the weakening inference can be omitted in P'.

Now consider the case of an $\wedge$ inference in P:

$$\frac{A, \Gamma_1 \qquad B, \Gamma_2}{A \wedge B, \Gamma_1, \Gamma_2}$$

For the first subcase, suppose that $A \wedge B$ is in category (α) or (γ), so $*(A \wedge B)$ is just $A \wedge B$. In this case, A and B are both in category (γ), so also $*(A) = A$ and $*(B) = B$. The $*$-translation of the $\wedge$ inference thus becomes

$$\frac{A, \Lambda, A \wedge B \qquad B, \Lambda, A \wedge B}{A \wedge B, \Lambda}$$

for suitable Λ, and this is still a valid inference. (The formula $A \wedge B$ appears in the upper cedents since the $*$-translations of the cedents A, Γ_1 and B, Γ_1 must contain $*(A \wedge B) = A \wedge B$.)

As the second subcase, suppose $A \wedge B$ is category (β), and thus A and B are also category (β). Expressing the formula B as a conjunction of its $\wedge$-components yields $B = B_1 \wedge B_2 \wedge \cdots \wedge B_k$ for $k \geq 1$. Let the pending implicants of $A \wedge B$ be $\overline{C}_1, \ldots, \overline{C}_\ell$ with $\ell \geq 0$. The formula B has the same pending implicants as $A \wedge B$. Similarly, $*(A)$ is $\overline{B}_1, \ldots, \overline{B}_k, \overline{C}_1, \ldots, \overline{C}_\ell$. Thus the $*$-translations of the cedents in the $\wedge$ inference become

$$\frac{\overline{B}_1, \ldots, \overline{B}_k, \overline{C}_1, \ldots, \overline{C}_\ell, \Lambda \qquad \overline{C}_1, \ldots, \overline{C}_\ell, \Lambda}{\overline{C}_1, \ldots, \overline{C}_\ell, \Lambda}$$

for suitable Λ. The dashed line is used to indicate that this is no longer a valid inference. However, since the lower cedent is the same as the upper right cedent, this inference can be completely omitted in P'.

Next consider the case of a cut inference in P:

$$\frac{A, \Gamma_1 \qquad \overline{A}, \Gamma_2}{\Gamma_1, \Gamma_2}$$

Clearly, A is of category (α), and $\overline{A}$ is of category (β). Since $\overline{A}$ has no pending implicants, $*(\overline{A})$ is the empty cedent; thus the $*$-translation of the three cedents has the form

$$\frac{*(A), \Lambda \qquad \Lambda}{\Lambda}$$

The cut inference therefore can be completely omitted in P'.

Now consider the case of an $\vee$ inference in P:

$$\frac{A, B, \Gamma}{A \vee B, \Gamma}$$

There are three subcases to consider. First, if $A \vee B$ is in category (γ), then so are A and B. The $*$-translation of the two cedents has the form

$$\frac{A, B, \Lambda, A \vee B}{A \vee B, \Lambda} \tag{5}$$

This of course is a valid inference, and remains in this form in P'.

The second subcase is when $A \vee B$ is category (α). Expressing A and B as disjunctions of their $\vee$-components yields $A = A_1 \vee \cdots \vee A_k$ and $B = B_1 \vee \cdots \vee B_\ell$ with $k, \ell \geq 1$. The $*$-translation of the $\vee$ inference is

$$\frac{A_1, \ldots, A_k, B_1, \ldots, B_k, \Lambda}{A_1, \ldots, A_k, B_1, \ldots, B_k, \Lambda}$$

and so this inference can be omitted in P'.

The third subcase is when $A \vee B$ is category (β). In this subcase, A and B are both category (γ). We have $*(A) = A$ and $*(B) = B$. And, $*(A \vee B)$ is $\overline{C}_1, \ldots, \overline{C}_k$, where the $\overline{C}_i$'s are the pending implicants of $A \vee B$, with $k \geq 0$. Thus, the $*$-translation of the cedents in the $\vee$ inference has the form

$$\frac{A, B, \Lambda, \overline{C}_1, \ldots, \overline{C}_k}{\overline{C}_1, \ldots, \overline{C}_k, \Lambda}$$

Of course, this is not a valid inference. Note that the formulas $\overline{C}_i$ must be included in the upper sequent since they are part of $*(A \vee B)$. From Lemma 4, the upper left sequent of the jump target of $A \vee B$ has $*$-translation of the form

$$\overline{A \vee B}, \overline{C}_1, \ldots, \overline{C}_k, \Lambda',$$

where $\Lambda' \subseteq \Lambda$. The following inferences are used in P' to replace the $\vee$ inference:

$$\text{Cut: } \dfrac{\overline{A \vee B}, \overline{C}_1, \ldots, \overline{C}_k, \Lambda' \qquad \dfrac{A, B, \Lambda, \overline{C}_1, \ldots, \overline{C}_k}{A \vee B, \Lambda, \overline{C}_1, \ldots, \overline{C}_k}}{\overline{C}_1, \ldots, \overline{C}_k, \Lambda} \tag{6}$$

This cut is permitted in P' since $A \vee B$ is an $\wedge$-component of a right cut formula in P. Note that the $\vee$ inference in P has been replaced in P' with two inferences, namely a cut and an $\vee$ inference.

Now consider the case of a $\forall$ inference in P

$$\dfrac{A(b), \Gamma}{(\forall x)A(x), \Gamma}$$

This case is handled similarly to the case of an $\vee$ inference. The formula $A(b)$ is category (γ), so $*(A(b)) = A(b)$. If the formula $(\forall x)A(x)$ is category (α) or (γ), then $*((\forall x)A(x)) = (\forall x)A(x)$. In this case, the $*$-translation of the $\forall$ inference gives

$$\dfrac{A(b), \Lambda, (\forall x)A(x)}{(\forall x)A(x), \Lambda}$$

for suitable Λ. This is still a valid inference, and is used as is in P'. Suppose, on the other hand, that $(\forall x)A(x)$ is category (β). In this case, the $*$-translation of the $\forall$ inference has the form

$$\begin{array}{c} A(b), \Lambda, \overline{C}_1, \ldots, \overline{C}_k \\ \hdashline \overline{C}_1, \ldots, \overline{C}_k, \Lambda \end{array}$$

where $\overline{C}_1, \ldots, \overline{C}_k$ are the pending implicants of $(\forall x)A(x)$. Note this is not a valid inference. By Lemma 4, the $*$-translation of the upper left cedent of the jump target of $(\forall x)A(x)$ is equal to

$$(\exists x)\overline{A(x)}, \overline{C}_1, \ldots, \overline{C}_k, \Lambda',$$

where $\Lambda' \subseteq \Lambda$. The following inferences are used in P' to replace the $\forall$ inference:

$$\text{Cut: } \dfrac{(\exists x)\overline{A(x)}, \overline{C}_1, \ldots, \overline{C}_k, \Lambda' \qquad \dfrac{A(b), \Lambda, \overline{C}_1, \ldots, \overline{C}_k}{(\forall x)A(x), \Lambda, \overline{C}_1, \ldots, \overline{C}_k}}{\overline{C}_1, \ldots, \overline{C}_k, \Lambda} \tag{7}$$

Note that since P is in free variable normal form, the variable b does not appear in the lower cedent of the new $\forall$ inference. The $\forall$ inference in P has been replaced in P' with two inferences: a cut and a $\forall$ inference.

The case of an $\exists$ inference in P is handled in exactly the same way as a $\forall$ inference. We omit the details.

The above completes the construction of P' from P. By construction, the inferences in P' are valid. To verify that P' is globally a valid proof, we need to ensure that it is acyclic, so there is no chain of inferences that forms a cycle. This follows immediately from the fact that the inferences in P' respect the post-order traversal of P. In particular, the upper left cedent of the jump target of a formula B comes before the cedent containing B in the post-order traversal of P. Therefore, P' is well-founded.

It is clear that P' can be constructed in polynomial time from P. The size of P' can be bounded as follows. First, each initial sequent in P can add at most one cut inference to P'. Each $\land$ inference in P can become at most one $\land$ inference in P'. Each $\lor$, $\forall$, and $\exists$ inference in P can become up to two inferences in P'. Each cut in P is replaced, at least locally, by zero inferences in P'. Let n_{Ax}, n_{Cut}, $n_\land$, $n_\lor$, $n_\forall$, and $n_\exists$ denote the numbers of initial sequents, cuts, $\land$, $\lor$, $\forall$, and $\exists$ inferences in P. Then $|P|$ equals $n_{\mathrm{Cut}}+n_\land+n_\lor+n_\forall+n_\exists$, and $|P'|$ is bounded by $n_{\mathrm{Ax}}+n_\land+2(n_\lor+n_\forall+n_\exists)$. Since w.l.o.g. there is at least one cut in P and since $n_{\mathrm{Ax}} = n_{\mathrm{Cut}} + n_\land + 1$, it follows that $|P'| \le 2 \cdot |P|$. Q.E.D. Theorem 3. $\Box$

4 Eliminating like quantifiers

We next show how to eliminate the outermost block of quantifiers from cut formulas.

Definition An $\exists$-*subformula* (resp., $\forall$-*subformula*) of A is a subformula that is contained in the scope of only $\exists$ (resp., $\forall$) quantifiers. An $\exists$-*component* (resp., $\forall$-*component*) of A is a minimal $\exists$- or $\forall$-subformula (respectively). A $\forall/\exists$-*component of a cut formula in* P is a $\forall$-component of a left cut formula in P or an $\exists$-component of a right cut formula in P.

Theorem 5 *Let P be a tree-like Tait calculus proof of Γ. Then there is a dag-like proof P', also of Γ, such that each cut formula of P' is a $\forall/\exists$-component of a cut formula of P, and such that $|P'| \le 4^{|P|/5} \le (1.32)^{|P|}$ and $h(P') \le |P|$. As a consequence of the height bound, P' can also be expressed as a tree-like proof of size $\le 2^{|P|}$. Similarly, $h(P') \le 2^{h(P)}$.*

Without loss of generality, P is in free variable normal form and satisfies the auxiliary condition. Each formula B in P can be put in one of the following categories (α)-(γ):

(α) B has a left cut formula A as a descendant and corresponds to a $\forall$-subformula of A, or

(β) B has a right cut formula A as a descendant and corresponds to an $\exists$-subformula of A, or

(γ) Neither (α) nor (β) holds.

Definition An $\exists$ inference as shown in Figure 1 is *critical* if the auxiliary formula $A(s)$ does not have an $\exists$ as its outermost connective. The formula $A(s)$ is also referred to as $\exists$-*critical.* If $A(s)$ is furthermore of category (β), then the $\exists$-*jump target* of $A(s)$ is the cut inference which has a descendant of $A(s)$ as a (right) cut formula. The $\exists$-*jump target cedent* of $A(s)$ is the upper left cedent of the jump target of $A(s)$. This is also referred to as the $\exists$-jump target cedent of the cedent Δ containing $A(s)$.

We now come to the crucial new definition for handling cut elimination of outermost like quantifiers. The intuition is that we want to trace, through the proof P, a possible branch in the proof P'. Along with this traced out path, we also need to keep a partial substitution assigning terms to variables: this substitution will track the needed term substitution for forming the corresponding cedent in P'. First we define an "$\exists$-path" and then we define the associated substitution.

Definition A cut inference is called *to-be-eliminated* if the outermost connective of the cut formula is a quantifier. An $\exists$-*path* π through P consists of a sequence of cedents Δ_1, Δ_2,..., Δ_m from P such that Δ_1 is the endsequent of P and such that for each $i < m$, one of the following holds:

- Δ_i is the lower cedent of a to-be-eliminated cut inference, and Δ_{i+1} is its right upper cedent, or
- Δ_i is the lower cedent of an inference other than a to-be-eliminated cut, and Δ_{i+1} is an upper cedent of the same inference, or
- Δ_i is the upper cedent of an $\exists$-critical inference, and Δ_{i+1} is the $\exists$-jump target cedent of Δ_i.

The $\exists$-path is said to *lead to* Δ_m.

It is easy to verify that, for Δ_i in π, the $\exists$-path π contains every cedent in P below Δ_i.

The cedents in an $\exists$-path are in reverse post-order from P. The effect of an $\exists$-path is to repeatedly traverse up to an $\exists$-critical inference — always going rightward at to-be-eliminated cuts — and then jump back down to the associated $\exists$-jump target cedent. The most important information needed to specify the $\exists$-path is the subsequence of cedents $\Delta_{i_1}, \Delta_{i_2}, \ldots, \Delta_{i_k}$, $i_1 < i_2 < \cdots < i_k$ which are $\exists$-critical and for which $\Delta_{i_\ell+1}$ is the $\exists$-jump target cedent of Δ_{i_ℓ}. The entire $\exists$-path can be uniquely reconstructed from this subsequence plus knowledge of the last cedent Δ_m in π.

There is a substitution σ_π associated with the $\exists$-path $\pi = \langle \Delta_1, \ldots, \Delta_m \rangle$. The domain of σ_π is the set of free variables appearing in or below Δ_m *plus* the set of outermost universally quantified variables occurring in the category (α) formulas in Δ_m.

Definition The definition of σ_π is by induction on the length of π. First, let $(\forall x_i)\cdots(\forall x_\ell)A$ be a formula in Δ_m in category (α) with $i \leq \ell$ such that A does not have outermost connective $\forall$. Since it is in category (α), this formula has the form $(\forall x_i)\cdots(\forall x_\ell)A(b_1, \ldots, b_{i-1}, x_i, \ldots, x_\ell)$, and has a descendant of the form $(\forall x_1)\cdots(\forall x_\ell)A(x_1, \ldots, x_\ell)$ which is the left cut formula of a cut inference. Since the cut is to-be-eliminated, π must reach the upper left cedent by way of a "jump" from an $\exists$-critical cedent $\Delta_i \in \pi$. As pictured, the associated $\exists$-critical formula must have the form $\overline{A(s_1, \ldots, s_\ell)}$:

$$\dfrac{\begin{matrix}(\forall x_i)\cdots(\forall x_\ell)A(b_1, \ldots, b_{i-1}, x_i, \ldots, x_\ell), \Gamma \\ \ddots\vdots\;\cdot^{\cdot^{\cdot}} \\ (\forall x_1)\cdots(\forall x_\ell)A(x_1, \ldots, x_\ell), \Gamma_1\end{matrix} \qquad \begin{matrix}\dfrac{\overline{A(s_1, \ldots, s_\ell)}, \Gamma'}{(\exists x_\ell)\overline{A(s_1, \ldots, s_{\ell-1}, x_\ell)}, \Gamma'} \\ \ddots\vdots\;\cdot^{\cdot^{\cdot}} \\ (\exists x_1)\cdots(\exists x_\ell)\overline{A(x_1, \ldots, x_\ell)}, \Gamma_2\end{matrix}}{\Gamma_1, \Gamma_2}$$

Note that the terms $s_1, \ldots, s_\ell$ are uniquely determined by π, since they are found by following the path from the upper right cedent of the cut inference to the cedent Δ_i, and setting the s_i's to be the terms used for $\exists$ inferences acting on the descendants of $A(\vec{s})$.

Let π' be π truncated to end at $\overline{A(\vec{s}), \Gamma}$. The substitution σ_π is defined to map the bound variables $x_i, \ldots, x_\ell$ to the terms $s_i\sigma_{\pi'}, \ldots, s_\ell\sigma_{\pi'}$. (Strictly speaking, the substitution σ_π acts on the occurrences of variables, since the same variable may be used in multiple quantifiers and in different formulas; this is suppressed in the notation however.)

For b a free variable appearing in or below Δ_m, the value $\sigma_\pi(b)$ is defined as follows. If there is a $\forall$ inference, below Δ_m,

$$\frac{A(b), \Gamma}{(\forall x)A(x), \Gamma}$$

that uses b as an eigenvariable, and if $(\forall x)A(x)$ is category (α), then define $\sigma_\pi(b)$ to equal the value of $\sigma_{\pi'}(a)$, where π' is π truncated to end at the lower cedent of the $\forall$ inference. For example, in the proof displayed above, $\sigma_\pi(b_i) = s_i$. Otherwise, if there is no such $\forall$ inference, define $\sigma_\pi(b) = b$.

Definition Let A be a formula appearing in a cedent Δ of P. Let π be an $\exists$-path leading to Δ. Then $*_\pi(A)$ is defined as follows:

- If A is in category (α) and has the form $A = (\forall x_1) \cdots (\forall x_\ell)B$ with $\ell > 0$ and B not starting with a $\forall$ quantifier, then define $*_\pi(A)$ to be the formula $B\sigma_\pi$, namely the formula obtained by replacing each x_i with $\sigma_\pi(x_i)$ and each free variable b with $\sigma_\pi(b)$.

- If A is in category (β) and has outermost connective $\exists$, then $*_\pi(A)$ is the empty cedent.

- Otherwise $*_\pi(A)$ is the formula $A\sigma_\pi$, namely obtained by replacing each free variable b with $\sigma_\pi(b)$.

For A appearing below Δ, we define $*_\pi(A)$ to equal $*_{\pi'}(A)$ where π' is π truncated to end at the cedent Δ' containing A. The $*_\pi$-*translation*, $*_\pi(\Delta)$, of Δ is the cedent containing exactly the formulas $*_\pi(A)$ for A appearing in or below Δ in P.

We can now give the proof of Theorem 5. The proof P' will be formed from the cedents $*_\pi(\Delta)$ where Δ ranges over the cedents of P, and π ranges over the $\exists$-paths leading to Δ. The inferences in P' will respect the post-ordering of P, and P' will be a dag.

As before, we must show how to connect up the cedents $*_\pi(\Delta)$ to make P' into a valid proof. The argument again splits into cases based on the type of inference used to infer Δ in P. The cases of initial cedents, $\vee$ inferences, $\wedge$ inferences, and weakenings are all immediate. These inferences remain valid after their cedents are replaced by their $*_\pi$-translations, since initial cedents contain only atomic formulas, and since the $*_\pi$-translations respect propositional connectives.

Consider the case where Δ is inferred by a $\forall$ inference in P:

$$\frac{A(b), \Gamma}{(\forall x)A(x), \Gamma}$$

The $\exists$-path π ends at the lower cedent Δ. Define π' to be the $\exists$-path that extends π by one step to the upper cedent Δ'. If $(\forall x)A(x)$ and $A(b)$ are not in category (α), then $\sigma_{\pi'}(b) = b$ and the inference is still valid since the $*_{\pi'}$-/$*_{\pi}$-translations of $A(b)$ and $(\forall x)A(x)$ are equal to $C(b)$ and $(\forall x)C(x)$ for C defined by $C(b) = A(b)\sigma_\pi = A(b)\sigma_{\pi'}$. Thus, in this case, the result is still a valid $\forall$ inference. Otherwise, $A(b)$ and $(\forall x)A(x)$ are both in category (α). In this case, $*_\pi(A(b)) = *_\pi((\forall x)A(x))$; the $\forall$ inference has equal upper and lower cedents and is just omitted from P'.

Now consider the case where Δ is inferred in P with an $\exists$ inference:

$$\frac{A(s), \Gamma}{(\exists x)A(x), \Gamma}$$

Define π' as in the previous case. If $A(s)$ and $(\exists x)A(x)$ are not in category (β), then the $*_\pi$-translation leaves the quantifier on x untouched, and the $*_{\pi'}$-/$*_\pi$-translation of the inference is still a valid inference in P'. Otherwise, both formulas are in category (β). If $A(s)$ has an $\exists$ as its outmost connective, then $*_{\pi'}(A(s))$ and $*_\pi((\exists x)A(x))$ are both empty, and the $*_{\pi'}$- and $*_\pi$-translations (respectively) of the upper and lower cedents are identical, and the $\exists$ inference can be omitted in P'. If A does not have an $\exists$ as its outermost connective, then the $*_{\pi'}$-/$*_\pi$-translations of the cedents in the inference are

$$\frac{*_{\pi'}(A(s)), \Lambda}{\Lambda}$$

where Λ contains the formulas $*_\pi(B)$ for all formulas B, other than $A(s)$, which occur in or below Δ in P. The upper left cedent of the $\exists$-jump target of $A(s)$ has the form

$$\Gamma_1, (\forall x_1) \cdots (\forall x_\ell)\overline{A(x_1, \ldots, x_\ell)},$$

where $x = x_\ell$ and $A(s) = A(s_1, \ldots, s_\ell)$ with s corresponding to the term s_ℓ. Let π'' be the $\exists$-path that extends π' by the addition of this upper left cedent. The $*_{\pi''}$-translation of the upper left cedent has the form

$$\Lambda_1, \overline{A(s_1, \ldots, s_\ell)}\sigma_{\pi''}.$$

Here $\Lambda_1 \subseteq \Lambda$, and $\overline{A(s_1, \ldots, s_\ell)}\sigma_{\pi''}$ is the same as $\overline{*_{\pi'}(A(s))}$. Hence, a cut inference gives

$$\frac{\Lambda_1, \overline{A(s_1,\ldots,s_\ell)}\sigma_{\pi''} \qquad *_{\pi'}(A(s)), \Lambda}{\Lambda}$$

The $\exists$ inference in P is thus replaced with a cut inference in P', but on a formula of lower complexity than the cut in P.

Finally consider the case of a cut inference in P as shown in Figure 1 with left cut formula A and right cut formula $\overline{A}$. First suppose it is not a to-be-eliminated cut. Let π_1 and π_2 be the $\exists$-paths which extend π by one step to include the upper left or right cedent of the cut, respectively. Then $*_{\pi_1}(A)$ and $*_{\pi_2}(\overline{A})$ are complements of each other, and the cut remains valid in P'. Otherwise, the cut is to-be-eliminated, and π_2 is again a valid $\exists$-path. The right cut formula $\overline{A}$ is category (β) and has outermost connective $\exists$. Thus $*_{\pi_2}(\overline{A})$ is the empty cedent, so the $*_{\pi_2}$-translation of the right upper cedent and the $*_\pi$-translation of the lower cedent are identical. In this case, the cut can be removed completely from P'.

The above completes the construction of P'. The next lemma will be used to bound its size.

Lemma 6 *let Δ be a cedent in P. The number of $\exists$-paths π to Δ in P is $\le (1.32)^{|P|}$.*

Proof Recall that an $\exists$-path π to Δ can be uniquely characterized by its final cedent $\Delta_m = \Delta$ and its subsequence $\Delta_{i_1}, \ldots, \Delta_{i_k}$ of cedents which are $\exists$-critical and have $\Delta_{i_\ell+1}$ the $\exists$-jump target cedent of Δ_{i_ℓ}. We will bound the number N of ways to select the $\exists$-critical cedents in this subsequence. For this, we group the $\exists$-critical cedents of P according to their $\exists$-jump target. Let there be m many to-be-eliminated cut inferences in P, and suppose that the i-th such cut has n_i many $\exists$-critical cedents associated with it. The i-th cut also has at least one $\forall$ inference associated with it that introduces a $\forall$ quantifier in its left cut formula. Therefore $|P| \ge \sum_{i=1}^m (n_i + 2)$. Each $\exists$-path π can jump from at most one of the n_i $\exists$-critical cedents associated with the i-th cut. It follows that there are at most $\prod_{i=1}^m (n_i + 1)$ many $\exists$-paths; namely, there are at most $n_i + 1$ choices for which one, if any, of i-th cut's associated $\exists$-critical cedents are included in π.

To upper bound the value $N = \prod_{i=1}^m (n_i + 1)$, take the logarithm, and upper bound $\sum_{i=1}^m \ln(n_i + 1)$ subject to $\sum_{i=1}^m (n_i + 2) \le |P|$. For integer values of x, $(\ln x)/(x+1)$ is maximized at $x = 4$. Thus, $\ln N \le |P| \cdot (\ln 4)/5$; that is, $N \le |P| \cdot 4^{|P|/5} \le (1.32)^{|P|}$. □

The size bound of Theorem 5 follows immediately from the lemma. Namely, P' contains at most one cedent for each path to each cedent Δ in P, and thus $|P'| \leq |P| \cdot (1.32)^{|P|}$. The height bound $h(P') \leq |P|$ follows from the construction of P, since paths π traverse cedents of P in reverse postorder, and each $\wedge$, $\vee$, $\exists$, $\forall$, and cut inference along π contributes at most inference to P'. (Note that cuts contribute an inference only when used as a jump target.) Q.E.D. Theorem 5

The proof P' was constructed in a highly uniform way from P. Indeed, P' can be generated with a polynomial time algorithm f that operates as follows: f takes as input a string w of length $\leq |P|$ many bits, and outputs whether the string w is an index for a cedent Δ_w in P', and if so, f also outputs: (a) the cedent Δ_w with terms specified as dags, and (b) what kind of inference is used to derive Δ_w, and (c) the index w' or indices w', w'' of the cedent(s) from which Δ_w is inferred in P'. For (a), note that the cedent Δ_w can be written out in polynomial length only if terms are written as dags (that is, circuits) rather than as trees (that is, as formulas). This is because the iterated application of substitutions may cause the terms $\sigma_\pi(b)$ to be exponentially big when written out as formulas instead of as circuits. Also note that, although some inferences in P' become trivial and are omitted in P', we can avoid using Rep inferences in P' by the simple convention that indices w that would lead to Rep inferences are taken to not be valid indices. (An example of this would be a w encoding an $\exists$-path leading to a to-be-eliminated cut.)

This means of course that there is a polynomial space algorithm that lists out the proof P'.

5 Eliminating and/exists and or/forall blocks

This section gives an algorithm for eliminating outermost blocks of $\vee/\forall$ (equivalently, $\wedge/\exists$) connectives from cut formulas, where the $\vee$ and $\forall$ (resp., $\wedge$ and $\exists$) connectives can be arbitrarily interspersed.

Definition A subformula B of A is an $\vee\forall$-*subformula* of A if B is in the scope of only $\vee$ and $\forall$ connectives. The $\vee\forall$-*components* of A are the minimal $\vee\forall$-subformulas of A. The $\wedge\exists$-*subformulas* and $\wedge\exists$-*components* of A are defined similarly.

An $\vee\forall/\wedge\exists$-*component* of a cut formula in P is either an $\vee\forall$-component of a left cut formula of P or an $\wedge\exists$-component of a right cut formula of P.

Theorem 7 *Let P be a tree-like Tait calculus proof of Γ. Then there is a dag-like proof P', also of Γ, such that each cut formula of P' is an $\wedge\exists/\vee\forall$-component of a non-atomic cut formula of P, and such that $|P'| \leq 4^{|P|/5} \leq$*

$(1.32)^{|P|}$ *and* $h(P') \leq |P|$. *Consequently,* P' *can also be expressed as a tree-like proof of size* $\leq 2^{|P|}$.

Note that all cuts in P are simplified in P'. The atomic cuts in P are eliminated when forming P'. However, new cuts are added on $\wedge\exists/\vee\forall$-components of cuts in P, and some of these might be cuts on atomic formulas. If all cuts formulas in P are atomic, then P' is cut free.

W.l.o.g., P is in free variable normal form and satisfies the auxiliary condition. Each formula B in P can be put in one of the following categories (α)-(γ):

(α) B has a left cut formula A as a descendant and corresponds to an $\vee\forall$-subformula of A, or

(β) B has a right cut formula A as a descendant and corresponds to an $\wedge\exists$-subformula of A, or

(γ) Neither (α) nor (β) holds.

Definition The *jump target* of a category (β) formula B occurring in P is the first cut or $\wedge$ inference below the cedent containing B that has some descendant of B as the auxiliary formula D in its right upper cedent. The jump target will again be of the form (1). Its right auxiliary formula D has a unique subformula B' which corresponds to B. B' occurs only in the scope of $\exists$ connectives and $\wedge$ connectives, and only in the first argument of $\wedge$ connectives. (The last part holds since otherwise the jump target would be an $\wedge$ inference higher in the proof.) The *jump target cedent* is defined as before.

Suppose a category (β) formula B has descendant D as the right auxiliary formula of its jump target. Let the $\wedge\exists$-components of D be $D_m, \ldots, D_k$ in left-to-right order. The $\wedge\exists$-components of B in left-to-right order can be listed as $B_m, \ldots, B_\ell$, with each B_i corresponding to D_i, with $1 \leq m \leq \ell \leq k$. The formulas $\overline{D}_{\ell+1}, \ldots, \overline{D}_k$ are the *pending implicants* of B. The *pending quantifiers* of B are the quantifiers $(\exists x)$ which appear to the right of the subformula D_ℓ in D and are outermost connectives of $\wedge\exists$-subformulas of D. Let B' be the subformula of D that corresponds to B; the *current quantifiers* of B are the quantifiers $(\exists x)$ in D which contain B' in their scope.

The pending implicants of B will be used similarly as in the proof of Theorem 3, but first we need to define $\wedge\exists$-paths and substitutions σ_π similarly to the proof of Theorem 5. Now, σ_π must also map the pending quantifier variables to terms.

Definition An upper cedent Δ of an $\wedge$ or $\exists$ inference is *critical* if the auxiliary formula in Δ is either atomic or has outermost connective $\vee$ or $\forall$.

Definition A cut inference in P is *non-atomic* if its cut formulas are not atomic. An $\wedge\exists$-*path* π through P consists of a sequence $\Delta_1, \ldots, \Delta_m$ of cedents from P such that Δ_1 is the end cedent of P and such that, for each $i < m$, one of following holds:

- Δ_i is the lower cedent of non-atomic cut inference, and Δ_{i+1} is its right upper cedent, or
- Δ_i is the lower cedent of an inference other than a non-atomic cut, and Δ_{i+1} is an upper cedent of the same inference, or
- Δ_i is a critical upper cedent of an $\wedge$ or $\exists$ inference with auxiliary formula A, and Δ_{i+1} is the jump target cedent of A.

The next definition of σ_π is more difficult than in the proof of Theorem 5 because the substitution has to act also on the pending implicants of category (β).

Definition Let π be $\wedge\exists$-path as above. The domain of the substitution σ_π is: the free variables appearing in or below Δ_m, the variables of the $\vee\forall$-outermost quantifiers of each category (α) formula in Δ_m, and the variables of the pending quantifiers of each category (β) formula in Δ_m.[3] The definition of σ_π is defined by induction on the length of π. For π containing just the end cedent, σ_π is the identity mapping with domain the parameter variables of P. Otherwise, let π' be the initial part of π up through the next-to-last cedent Δ_{m-1} of π, and suppose $\sigma_{\pi'}$ is already defined. There are several cases to consider.

a. Suppose Δ_{m-1} and Δ_m are the lower cedent and an upper cedent of some inference other than a $\forall$ inference. The σ_π is same as $\sigma_{\pi'}$.

b. Suppose Δ_{m-1} and Δ_m are the lower cedent and an upper cedent of a $\forall$ inference as shown in Figure 1. If the principal formula $(\forall x)A(x)$ is category (α), then σ_π extends $\sigma_{\pi'}$ by letting $\sigma_\pi(b) = \sigma_{\pi'}(x)$ where $(\forall x)$ is the quantifier introduced by the $\forall$ inference. Otherwise, $\sigma_\pi(b) = b$. And, σ_π is equal to $\sigma_{\pi'}$ for all other variables in its domain.

[3]As before, strictly speaking, a variable might be quantified at multiple places, and σ acts on variables according to how they are bound by a quantifier, but we suppress this in the notation.

c. Otherwise, Δ_m is the jump target cedent of Δ_{m-1}. Suppose the jump target is an $\wedge$ inference

$$\frac{C, \Gamma_1 \qquad D, \Gamma_2}{C \wedge D, \Gamma_1, \Gamma_2}$$

For b a free variable in C, Γ_1, the value $\sigma_\pi(b)$ is defined to equal $\sigma_{\pi'}(b)$. Similarly, for any pending quantifier $(\exists x)$ of any category (β) formula in Γ_1 and for any $\vee\forall$-outermost quantifier $(\forall x)$ of any category (α) formula in Γ_1, set $\sigma_\pi(x) = \sigma_{\pi'}(x)$.

We also must define the action of σ_π on the pending quantifiers of the category (β) formula C. Let D_1 be the first (leftmost) $\wedge\exists$-component of D. The cedent Δ_{m-1} has the form B_1, Γ_3 where B_1 is an ancestor of D and corresponds to D_1. Write $D_1 = D_1(x_1, \ldots, x_j)$ where $(\exists x_1)$,..., $(\exists x_j)$ are the current quantifiers for D_1. Then $B_1 = B_1(s_1, \ldots, s_j)$ where the s_i's are the terms used for $\exists$ inferences acting on descendants of B_1. The $(\exists x_i)$'s are pending quantifiers of C, and $\sigma_\pi(x_i)$ is defined to equal $s_i\sigma_{\pi'}$. The rest of the pending quantifiers of C are the pending quantifiers of B_1 in the cedent Δ_{m-1}: for these variables, σ_π is defined to equal the value of $\sigma_{\pi'}$.

d. Suppose that Δ_m is the left upper cedent of the jump target of Δ_{m-1}, and the jump target is a cut inference

$$\frac{\overline{D}, \Gamma_1 \qquad D, \Gamma_2}{\Gamma_1, \Gamma_2}$$

Let π' be as before, and set $\sigma_\pi(b) = \sigma_{\pi'}(b)$ for all free variables of the lower cedent. For any pending quantifier $(\exists x)$ of any category (β) formula in Γ_1 and for any $\vee\forall$-outermost quantifier $(\forall x)$ of any category (α) formula in Γ_1, set $\sigma_\pi(x) = \sigma_{\pi'}(x)$. Now, let $D_1 = D_1(x_1, \ldots, x_j)$ and $B_1 = B_1(s_1, \ldots, s_j)$ as in the previous case. Consider any $\vee\forall$-outermost quantifier $(\forall y)$ of $\overline{D}$. If y is one of the x_i's, define $\sigma_\pi(y) = s_i\sigma_{\pi'}$. Otherwise, $(\exists y)$ is a pending quantifier of D_1, and a pending quantifier of B_1 in Δ_{m-1}, and we define $\sigma_\pi(y) = \sigma_{\pi'}(y)$.

Definition Suppose A is a formula occurring in cedent Δ in P, and π is an $\wedge\exists$-path leading to Δ. The formula $*_\pi(A)$ is defined as follows:

- If A is category (β), then $*_\pi(A)$ is the cedent containing the formulas $B\sigma_\pi$ for each pending implicant B of A.

- If A is category (α), then $*_\pi(A)$ is the cedent containing $B\sigma_\pi$ for each $\vee\forall$-component B of A.

- Otherwise $*_\pi(A)$ is $A\sigma_\pi$.

The notation $*_\pi(A)$ is extended to apply also to A appearing in a cedent Δ' below the cedent Δ. Let π' be the initial subsequence of π leading to Δ'. Then define $*_\pi(A) = *_{\pi'}(A)$. The $*_\pi$-translation of Δ consists of the formulas $*_\pi(A)$ such that A appears in or below Δ in P.

The next lemma is analogous to Lemma 4.

Lemma 8 *Suppose B is a category (β) formula in a cedent Δ in P, and let π be an $\wedge\exists$-path to Δ. Also suppose B does not have outermost connective $\wedge$ or $\exists$. Let $\overline{C}_1, \ldots, \overline{C}_m$ be the pending implicants of B. Let Δ' be B's jump target cedent, and E be the auxiliary formula in Δ'. Then there is an $\wedge\exists$-path π' to Δ' such that $*_{\pi'}(E)$ equals the cedent $\overline{B}\sigma_\pi, \overline{C}_1\sigma_\pi, \ldots, \overline{C}_m\sigma_\pi$.*

Proof The jump target of B is either a cut or an $\wedge$ inference as shown in (1), with B corresponding to the first $\wedge\exists$-component C_0 of D. The remaining $\wedge\exists$-components of D are $C_1,\ldots,C_r$ where $0 \le r \le m$. Of course, their negations are (some of the) pending implicants of B.

Suppose the jump target is a non-atomic cut inference. Then we have $r = m$. Since B does not have outermost connective $\wedge$ or $\exists$ and since the cut formula D is non-atomic, B is not the same as D. Consider the lowest direct descendant of B; it appears in a cedent Δ'', and is the auxiliary formula of an $\exists$ inference, or the left auxiliary formula of an $\wedge$ inference. In either case, Δ'' is critical. Let π'' be the $\wedge\exists$-path consisting of the initial part of π to Δ''. Set π' to be the $\wedge\exists$-path that follows π'' and then jumps from Δ'' to the upper left cedent Δ' of the jump target. The left cut formula E is equal to $\overline{D}$, and the $\vee\forall$-components of E are $\overline{C}_0, \ldots, \overline{C}_m$. The cedent $*_{\pi'}(E)$ consists of the formulas $\overline{C}_i\sigma_{\pi'}$. For $i = 0$, $\sigma_{\pi'}$ was defined so that $C_0\sigma_{\pi'} = B\sigma_\pi$. Likewise, for $i > 0$, we have $C_i\sigma_{\pi'} = C_i\sigma_{\pi''}$. Also, by cases a. and b. of the definition of σ_π, we have $C_i\sigma_{\pi''} = C_i\sigma_\pi$. Thus the lemma holds.

Second, suppose the jump target is a cut on an atomic formula. The right cut formula is equal to B of course; the left cut formula E is equal to $\overline{B}$. Letting π' be as above, $*_{\pi'}(E)$ is equal to $\overline{B}\sigma_{\pi'} = \overline{B}\sigma_\pi$ as desired.

Now suppose the jump target is an $\wedge$ inference, as in (1), where $E = C$. If D is atomic, then D is a direct descendant of B (possibly even the same occurrence as B). In this case, let Δ'' be the cedent containing D (the upper right cedent of the $\land$ inference), let π'' be the initial part of π' leading to Δ'', and let π' be π'' plus the upper left cedent Δ'. (Note that Δ' is the jump target cedent of D.) Then, the pending implicants of C in Δ' are $\overline{D} = \overline{B}$ and $\overline{C}_1, \ldots, \overline{C}_m$. We have $D\sigma_{\pi'} = D\sigma_{\pi''} = D\sigma_\pi$ and also $C_i\sigma_{\pi'} = C_i\sigma_{\pi''} = C_i\sigma_\pi$, so the lemma holds. Now suppose D is not atomic. Then define π', π'', and Δ'' exactly as in the case above where jump target of B was a cut inference. The pending implicants of C are $\overline{C}_0, \ldots, \overline{C}_m$, and, as before, we have $C_0\sigma_{\pi'} = \overline{B}\sigma_{\pi''} = B\sigma_\pi$ and $C_i\sigma_{\pi'} = C_i\sigma_{\pi''} = C_i\sigma_\pi$, satisfying the conditions of the lemma. □

Proof (of Theorem 7.) The proof combines the constructions from the proofs of the two previous theorems. For each cedent Δ in P and each $\land\exists$-path leading to Δ, form the cedent Δ^π as the $*_\pi$-translation of Δ. Our goal is to show that these cedents can be combined to form a valid proof P'. The proof splits into cases to handle the different kinds of inferences in P separately. In each case, we have a cedent Δ and an $\land\exists$-path π leading to Δ, and need to show how Δ^π is derived in P'.

For the first case, consider an initial cedent Δ of the form $B, \overline{B}$ in P. As the first subcase, suppose neither B nor $\overline{B}$ is category (β). Then Δ^π is the cedent $B\sigma_\pi, \overline{B}\sigma_\pi, \Lambda$ where Λ is the cedent of formulas $*_\pi(E)$ for E a formula appearing below Δ in P. This is obtained in P' by applying a weakening to the initial cedent $B\sigma_\pi, \overline{B}\sigma_\pi$.

For the second subcase, suppose B is category (β) and $\overline{B}$ is not. The formula B has a right cut formula as descendant, and corresponds to the ℓ-th $\land\exists$-component D_ℓ of D. Let the pending implicants of B be $\overline{D}_{\ell+1}, \ldots, \overline{D}_k$. By Lemma 8, there is an $\land\exists$-path π' to the upper left cedent Δ' of the jump target such that the auxiliary formula E in Δ' has $*_{\pi'}(E)$ equal to $\overline{B}\sigma_\pi, \overline{D}_\ell\sigma_\pi, \ldots, \overline{D}_k\sigma_\pi$. Thus, Δ^π and $(\Delta')^{\pi'}$ are

$$\overline{B}\sigma_\pi, \overline{D}_{\ell+1}\sigma_\pi, \ldots, \overline{D}_k\sigma_\pi, \Lambda$$

and

$$\overline{B}\sigma_\pi, \overline{D}_{\ell+1}\sigma_\pi, \ldots, \overline{D}_k\sigma_\pi, \Lambda'$$

where $\Lambda' \subseteq \Lambda$. In P', the first cedent is derived from the second by a weakening inference.

In the third subcase, both B and $\overline{B}$ are category (β). We have $*_\pi(B)$ still equal to $\overline{D}_{\ell+1}\sigma_\pi, \ldots, \overline{D}_k\sigma_\pi$, and now $*_\pi(\overline{B})$ is equal to $\overline{D}''_{\ell''+1}\sigma_\pi, \ldots, \overline{D}''_{k''}\sigma_\pi$ with the $\overline{D}''_i$'s the k'' pending implicants of $\overline{B}$. Using Lemma 8 twice, we have $\land\exists$-paths π' and π'' leading to cedents Δ' and Δ'' such that $(\Delta')^{\pi'}$ and

$(\Delta'')^{\pi''}$ (respectively) are

$$\overline{B}\sigma_\pi, \overline{D}_{\ell+1}, \ldots, \overline{D}_k\sigma_\pi, \Lambda'$$

and

$$B\sigma_\pi, \overline{D}''_{\ell''+1}, \ldots, \overline{D}''_{k''}\sigma_\pi, \Lambda''$$

where $\Lambda', \Lambda'' \subseteq \Lambda$. In P', using a cut and then a weakening gives Δ^π as desired.

Second, consider the (very simple) case where the cedent Δ is inferred by a weakening inference

$$\frac{\Delta'}{\Delta}$$

where $\Delta \subset \Delta'$. The path π to Δ can be extended by one more cedent to be a path π' to the cedent Δ'. The cedents Δ^π and $(\Delta')^{\pi'}$ are identical. Thus the weakening inference in P is just omitted in P'.

Now consider the case where Δ is the lower cedent of an $\wedge$ inference in P:

$$\frac{A, \Gamma_1 \qquad B, \Gamma_2}{A \wedge B, \Gamma_1, \Gamma_2}$$

Let Δ_1 and Δ_2 be the left and right upper cedents, respectively, and let π_1 and π_2 be the $\wedge\exists$-paths obtained by adding Δ_1 or Δ_2, respectively, to the end of π. First, suppose $A \wedge B$ is category (α) or (γ), so $*_\pi(A \wedge B)$ is $(A \wedge B)\sigma_\pi$. Then A and B are both category (γ), and $*_{\pi_1}(A) = A\sigma_{\pi_1} = A\sigma_\pi$ and $*_{\pi_2}(B) = B\sigma_{\pi_2} = B\sigma_\pi$. Thus, in P', the $\wedge$ inference becomes

$$\frac{A\sigma_\pi, \Lambda, (A \wedge B)\sigma_\pi \qquad B\sigma_\pi, \Lambda, (A \wedge B)\sigma_\pi}{(A \wedge B)\sigma_\pi, \Lambda}$$

and this is still a valid $\wedge$ inference.

For the second subcase, suppose $A \wedge B$, thus A and B, are category (β). The formula B in Δ_2 has the same pending implicants $\overline{C}_1, \ldots, \overline{C}_\ell$ as the formula $A \wedge B$ in Δ. Also, $\overline{C}_i\sigma_{\pi_2} = \overline{C}_i\sigma_\pi$. Thus Δ^π is the same as $(\Delta_2)^{\pi_2}$. This means that the $\wedge$ inference can be omitted in P'.

Next consider the case where Δ is the lower cedent of a cut in P:

$$\frac{A, \Gamma_1 \qquad \overline{A}, \Gamma_2}{\Gamma_1, \Gamma_2}$$

Let Δ_1 and Δ_2 be the left and right upper cedents, respectively, and π_1 and π_2 be the extensions of π to Δ_1 and Δ_2. The occurrence of $\overline{A}$ is category (β) of course, and $*_{\pi_2}(\overline{A})$ is the empty cedent. Thus, the cedents Δ^{π} and $(\Delta_2)^{\pi_2}$ are identical, and the cut inference may be omitted from P'.

Next consider the case where Δ is the lower cedent of an $\vee$ inference:

$$\frac{A, B, \Gamma}{A \vee B, \Gamma}$$

In this, and the remaining cases, let Δ' be the upper cedent of the inference, and let π' be π extended to the cedent Δ'. For the $\vee$ inference, $\sigma_{\pi'}$ is identical to σ_{π}. As a first subcase, suppose $A \vee B$ is category (γ), and thus A and B are as well. In this subcase, $*_{\pi}(A \vee B) = (A \vee B)\sigma_{\pi}$, $*_{\pi'}(A) = A\sigma_{\pi}$, and $*_{\pi'}(B) = B\sigma_{\pi}$. The $*_{\pi}$-translation of the two cedents thus forms a valid $\vee$ inference in P'.

The second subcase is when $A \vee B$, A, and B are category (α). Letting $A_1, \ldots, A_k$ be the $\vee\forall$-components of A, and $B_1, \ldots, B_{k'}$ be those of B, the $*_{\pi}$-translation of the $\vee$ inference has the form

$$\frac{A_1\sigma_{\pi}, \ldots, A_k\sigma_{\pi}, B_1\sigma_{\pi}, \ldots, B_{k'}\sigma_{\pi}, \Lambda}{A_1\sigma_{\pi}, \ldots, A_k\sigma_{\pi}, B_1\sigma_{\pi}, \ldots, B_{k'}\sigma_{\pi}, \Lambda}$$

and this can be omitted from P'.

The third subcase is when $A \vee B$ is category (β). Then A and B are category (γ), and $*(A) = A\sigma_{\pi}$ and $*(B) = B\sigma_{\pi}$. Also, $*(A \vee B)$ is $\overline{C}_1\sigma_{\pi}, \ldots, \overline{C}_k\sigma_{\pi}$, where the $\overline{C}_i$'s are the pending implicants of $A \vee B$. Thus, the $*_{\pi}$-translation of the cedents in the $\vee$ inference has the form

$$\frac{A\sigma_{\pi}, B\sigma_{\pi}, \Lambda, \overline{C}_1\sigma_{\pi}, \ldots, \overline{C}_k\sigma_{\pi}}{\overline{C}_1\sigma_{\pi}, \ldots, \overline{C}_k\sigma_{\pi}, \Lambda} \tag{8}$$

Of course, this is not a valid inference. Let Δ'' be the upper left cedent of the jump target of $A \vee B$. From Lemma 8, there is an $\wedge\exists$-path π'' leading to Δ'' so that the $*_{\pi''}$-translation of Δ'' is

$$(\overline{A \vee B})\sigma_{\pi}, \overline{C}_1\sigma_{\pi}, \ldots, \overline{C}_k\sigma_{\pi}, \Lambda'$$

where $\Lambda' \subseteq \Lambda$. In P', this cedent and the upper cedent of (8) are combined with an $\vee$ inference and a cut to yield the lower cedent of (8), similarly to what was done in (6).

Now consider the case where Δ is the lower cedent of a $\forall$ inference

$$\frac{A(b), \Gamma}{(\forall x)A(x), \Gamma}$$

First suppose $(\forall x)A(x)$ is category (γ), so $*_\pi((\forall x)A(x)) = (\forall x)A(x)\sigma_\pi = (\forall x)A(x)\sigma_{\pi'}$. The formula $A(b)$ is category (γ) and $\sigma_{\pi'}(b) = b$, thus $*_{\pi'}(A(b)) = A(b)\sigma_\pi$. The $\forall$ inference of P becomes

$$\frac{A(b)\sigma_\pi, \Lambda, (\forall x)A(x)\sigma_\pi}{(\forall x)A(x)\sigma_\pi, \Lambda}$$

and this forms a valid $\forall$ inference in P'.

For the second subcase, suppose that $(\forall x)A(x)$ is category (β). Hence, $A(b)$ is category (γ). This case is similar to the third subcase for $\vee$ inferences above. We have $*_\pi((\forall x)A(x))$ equal to $\overline{C}_1\sigma_\pi, \ldots, \overline{C}_k\sigma_\pi$ where the $\overline{C}_i$'s are the pending implicants of $(\forall x)A(x)$. And, $*_\pi(A(b))$ equals $A(b)\sigma_\pi$; note $\sigma_\pi(b) = b$. Thus, the $*_{\pi'}$-/$*_\pi$-translation of the cedents in the $\forall$ inference has the form

$$\frac{A(b)\sigma_\pi, \Lambda, \overline{C}_1\sigma_\pi, \ldots, \overline{C}_k\sigma_\pi}{\Lambda, \overline{C}_1\sigma_\pi, \ldots, \overline{C}_k\sigma_\pi} \tag{9}$$

which is not a valid inference. Let Δ'' be the upper left cedent of the jump target of $(\forall x)A(x)$. By Lemma 8, there is an $\wedge\exists$-path π'' leading to Δ'' so that the $*_{\pi''}$-translation of Δ'' is

$$\overline{(\forall x)A(x)}\sigma_\pi, \overline{C}_1\sigma_\pi, \ldots, \overline{C}_k\sigma_\pi, \Lambda'$$

where $\Lambda' \subseteq \Lambda$. In P', this cedent and the upper cedent of (9) are combined with an $\forall$ inference and a cut to yield the lower cedent of (9), similarly to what was done in (7).

For the third subcase, suppose that $(\forall x)A(x)$ is category (α), so $A(b)$ is also category (α). By definition, $\sigma_{\pi'}(b) = \sigma_\pi(x)$. Thus, $*_{\pi'}(A(b)) = A(b)\sigma_{\pi'} = A(x)\sigma_\pi$. Also, $*_\pi((\forall x)A(x)) = A(x)\sigma_\pi$. Therefore, in P', the $\forall$ inference becomes trivial with Δ^π and $(\Delta')^{\pi'}$ equal to each other; so this inference is omitted from P'.

Finally, consider the case where Δ is the lower cedent of an $\exists$ inference

$$\frac{A(s), \Gamma}{(\exists x)A(x), \Gamma}$$

Note that $\sigma_{\pi'}$ is the same as σ_π. For the first subcase, suppose $(\exists x)A(x)$ is either category (α) or (γ), so $A(s)$ is category (γ). This gives $*_{\pi'}(A(s)) = A(s)\sigma_{\pi'} = A(s)\sigma_\pi$. And, since its outermost connective is $\exists$, $*_\pi((\exists x)A(x)) = (\exists x)A(x)\sigma_\pi$. The $\exists$ inference in P becomes, in P',

$$\frac{A(s)\sigma_\pi, \Lambda, (\exists x)A(x)\sigma_\pi}{(\exists x)A(x)\sigma_\pi, \Lambda}$$

which is a valid $\exists$ inference.

For the second subcase, suppose $(\exists x)A(x)$ and hence $A(s)$ are category (β). The two formulas have the same pending implicants, $\overline{C}_1, \ldots, \overline{C}_k$, for $k \geq 0$. Thus, $*_{\pi'}(A(s))$ and $*_\pi((\exists x)A(x))$ are both equal to the cedent $\overline{C}_1\sigma_\pi, \ldots, \overline{C}_k\sigma_\pi$. That is to say, Δ^π and $(\Delta')^{\pi'}$ are identical, and thus the $\exists$ inference can be omitted from P'.

The above completes the construction of P' from P. The discussion at the end of the proof of Theorem 5 applies equally well to the P' just constructed, and P' is again polynomial time uniform. □

6 Bounds on eliminating all cuts

This section gives bounds on eliminating all cuts from a proof. The bound obtained has the form $2^{|P|}_{d+O(1)}$, where d is the maximum quantifier alternation of cut formulas in P. The first-order formula classes Σ_i and Π_i are defined as usual by counting alternations of quantifiers, allowing propositional connectives to appear arbitrarily. Namely, $\Sigma_0 = \Pi_0$ is the set of quantifier free formulas; and, using Bachus-Naur notation,

$$\Sigma_i \quad ::= \quad \Sigma_{i-1} | \Pi_{i-1} | \Sigma_i \wedge \Sigma_i | \Sigma_i \vee \Sigma_i | \neg \Pi_i | (\exists x)\Sigma_i$$

$$\Pi_i \quad ::= \quad \Pi_{i-1} | \Sigma_{i-1} | \Pi_i \wedge \Pi_i | \Pi_i \vee \Pi_i | \neg \Sigma_i | (\forall x)\Pi_i$$

The *alternating quantifier depth (aqd)* of a cut is the minimum $i > 0$ such that one cut formula is in Σ_i and the other is in Π_i. The *alternation quantifier depth* of a proof P, denoted $\mathsf{aqd}(P)$, is the maximum aqd of any cut in P.

Theorem 9 *Let P be a tree-like proof, and let $d = \mathsf{aqd}(P)$. There is a cut free proof P' with the same end cedent as P with the size of P' bounded by $|P'| \leq 2^{|P|}_{d+O(1)}$.*

The proof of the theorem depends only on Theorem 5, not on Theorems 3 and 7. We also use upper bounds on eliminating cuts on quantifier free formulas as can be found in [17, 7, 5].

Proof It is helpful to briefly review the well-known fact that the size of formulas appearing in the tree-like proof P can be bounded by the number of inferences in P plus the size of the formulas in the end cedent of P. For this, recall that any formula B appearing in P has a unique descendent A such that A either is a cut formula or is in the end cedent of P. In addition,

B corresponds to a unique subformula C of A. Let C be a non-atomic subformula of a formula D in P which has a cut formula as descendant. If there is some ancestor B of D such that B corresponds to C and such that B is a principal formula of a logical inference, then leave C unchanged. If there is no such ancestor D, then mark C for deletion. Now replace every maximal subformula C in P marked for deletion with with an arbitrary atomic formula, say with $d{=}d$ for d some new free variable. The proof remains a valid proof (since only atomic formulas are allowed in initial cedents), and its end cedent is unchanged. Clearly, in the resulting proof, every cut formula has number of logical connectives bounded by the total number of $\land$, $\lor$, $\exists$ and $\forall$ inferences in P. Without loss of generality, we assume this is true of the proof P itself.

The main step in proving Theorem 9 is to convert P into a proof in which all cuts are in prenex form. As a preliminary step, we show that we may assume w.l.o.g. that no cut formula in P has multiple quantifiers on the same bound variable, or in other words, that the bound variables in a cut formula are distinct. Towards this end, for each cut inference in P, with formulas A and $\overline{A}$ as its cut formulas, rename the bound variables in A so that the quantifiers in A use distinct bound variables. This also renames the bound variables of $\overline{A}$ of course. Furthermore, if B is a formula with descendent A or $\overline{A}$, this induces a renaming of the bound variables in B according to the renaming of bound variables in the subformula of A or $\overline{A}$ that corresponds to B. By applying these renamings to all such formulas B, and repeating for all cuts in P, we obtain a proof with the same end cedent as P such that bound variables are never reused in cut formulas.[4] So, we may assume w.l.o.g. that P satisfies this property.

Now, for each cut in P, with cut formulas A and $\overline{A}$, choose an arbitrary prenex form A' for A so that the aqd of A' is $\leq \mathsf{aqd}(P)$. The formula A' is obtained by choosing an ordering of the quantifiers in A which respects the scope of the quantifiers, and then using standard prenex operations to move the quantifiers out to the front of the formula in the chosen order. The prenex form $(\overline{A})'$ of $\overline{A}$ is chosen with the same ordering and thus equals $\overline{A'}$.

Let B be any formula in P with a cut formula A as descendent. The quantifiers of A are ordered as just discussed to form its prenex form A'. Since B corresponds to a subformula of A, this induces an ordering on the quantifiers of B; the prenex form B' of B is defined using this induced ordering. On the other hand, if B has a descendent in the end cedent of P,

[4] The same construction could also rename bound variables in the end cedent of P, but this would then change the end cedent.

the formula B' is defined to be equal to B. For any cedent Δ in P, define Δ' to contain exactly the formulas B' for $B \in \Delta$.

The proof P' will contain the cedents Δ' for all $\Delta \in P$. However, the $\wedge$ and $\vee$ inferences in P may no longer be valid in P'. Cuts, weakenings, and quantifier inferences of P do remain valid in P'. In addition, since only atomic formulas are allowed initial cedents, the initial cedents of P are unchanged in P'.

In order to make P' a valid proof, we must replace the $\wedge$ and $\vee$ inferences of P with some new subproofs and cuts. The next lemma gives the key construction needed for this.

Lemma 10 *Let $B \wedge C$ be the principal formula of an $\wedge$ inference in P with a cut formula as descendent. The auxiliary formulas of the inference are B and C. Let B', C', and $(B \wedge C)'$ be their prenex forms in P'. Then the cedent*

$$\overline{B'}, \overline{C'}, (B \wedge C)' \tag{10}$$

has a cut free proof of length linear in the lengths of B and C. Similarly, if $B \vee C$ is the principal formula of an $\vee$ inference of P, then the cedents

$$\overline{B'}, (B \vee C)' \qquad \text{and} \qquad \overline{C'}, (B \vee C)' \tag{11}$$

have cut free proofs of length linear in the lengths of B and C.

Proof Let B' and C' have the forms $\mathcal{Q}_1 B_0$ and $\mathcal{Q}_2 C_0$ where $\mathcal{Q}_1$ and $\mathcal{Q}_2$ denote blocks of zero or more quantifiers and where B_0 and C_0 are quantifier free. The formula $(B \wedge C)'$ or $(B \vee C)'$ will have the form $\mathcal{Q}(B_0 \wedge C_0)$ or $\mathcal{Q}(B_0 \vee C_0)$. Here the quantifier block $\mathcal{Q}$ is obtained by arbitrarily interleaving (or, "shuffling") the two blocks $\mathcal{Q}_1$ and $\mathcal{Q}_2$.

We claim that, for any quantifier blocks $\mathcal{Q}_1$ and $\mathcal{Q}_2$, and any block $\mathcal{Q}$ obtained as a shuffle of $\mathcal{Q}_1$ and $\mathcal{Q}_2$, the cedents (10) and (11) have tree-like, cut free proofs with size equal to the number of logical connectives in the cedents being proved. This is proved by induction on the number of quantifiers in $\mathcal{Q}$.

The base case of the induction is when $\mathcal{Q}$ is empty, and B and C are quantifier free. As is well known (and easy to verify) there are proofs of the cedents $\overline{B}_0, B_0$ and $\overline{C}_0, C_0$ with sizes equal to twice the number of logical connectives in B_0 and C_0, respectively. These two cedents plus a single $\wedge$ or $\vee$ inference suffices to derive any of the cedents in (10) or (11).

For the induction step, suppose that $\mathcal{Q}$ contains at least one quantifier. The first quantifier can have the form $(\exists x)$ or $(\forall x)$ and is also the first quantifier of either $\mathcal{Q}_1$ or $\mathcal{Q}_2$. For instance, suppose $(\exists x)$ is the outermost quantifier of $\mathcal{Q}$ and $\mathcal{Q}_1$. Writing $B_0 = B_0(x)$ to show the occurrences of the bound variable x, and replacing occurrences of x with a new free variable a, the induction hypothesis gives derivations of the cedents

$$\overline{\mathcal{Q}_1^- B_0(a)}, \mathcal{Q}^-(B_0(a) \vee C_0) \quad \text{and} \quad \overline{\mathcal{Q}_2 C_0}, \mathcal{Q}^-(B_0(a) \vee C_0)$$

or

$$\overline{\mathcal{Q}_1^- B_0(a)}, \overline{\mathcal{Q}_2 C_0}, \mathcal{Q}^-(B_0(a) \wedge C_0)$$

where $\mathcal{Q}_1^-$ and $\mathcal{Q}^-$ are the blocks $\mathcal{Q}_1$ and $\mathcal{Q}$ minus the first quantifier $\exists x$. For the $\vee$ case, the derivation

$$\frac{\overline{\mathcal{Q}_2 C_0}, \mathcal{Q}^-(B_0(a) \vee C_0)}{\overline{\mathcal{Q}_2 C_0}, (\exists x)\mathcal{Q}^-(B_0(x) \vee C_0)}$$

gives the desired derivation of $\overline{\mathcal{Q}_2 C_0}, \mathcal{Q}(B_0 \vee C_0)$; and the derivation

$$\cfrac{\cfrac{\overline{\mathcal{Q}_1^- B_0(a)}, \mathcal{Q}^-(B_0(a) \vee C_0)}{\overline{\mathcal{Q}_1^- B_0(a)}, (\exists x)\mathcal{Q}^-(B_0(x) \vee C_0)}}{(\forall x)\overline{\mathcal{Q}_1^- B_0(x)}, (\exists x)\mathcal{Q}^-(B_0(x) \vee C_0)}$$

gives the desired derivation of $\overline{\mathcal{Q}_1 B_0}, \mathcal{Q}(B_0 \vee C_0)$. Note that the second inference is a $\forall$ inference; by the assumption of distinctness of bound variables, the eigenvariable a does not appear in C_0.

A similar argument works for the $\wedge$ case. The cases where outermost quantifier of $\mathcal{Q}$ is $(\forall x)$ are also similar. □

We can now complete the proof of Theorem 9. The proof P' is formed from the cedents Δ' defined above. Using the cedents Δ' maintains the validity of all inferences except for some of the $\vee$ and $\wedge$ inferences. In P' these inferences become

$$\frac{B', \Gamma_1' \qquad C', \Gamma_2'}{(B \wedge C)', \Gamma_1, \Gamma_2} \qquad \text{and} \qquad \frac{B', C', \Gamma}{(B \vee C)', \Gamma}$$

and these are no longer valid if their principal formula contains quantifiers and has a cut formula as descendent. However, the $\wedge$ inference can be simulated by using two cuts against the cedent $\overline{B'}, \overline{C'}, (B \wedge C)'$ given by Lemma 10. Likewise, the $\vee$ inference can be simulated by using two cuts with the cedents $\overline{B'}, (B \vee C)'$ and $\overline{C'}, (B \vee C)'$. This process replaces one inference in P with two cuts in P'; in addition, P' must contain the derivations of the cedents as given by Lemma 10. Since the formulas $(B \wedge C)'$ and $(B \vee C)'$ have cut formulas as descendents, their sizes are bounded by $|P|$ as discussed at the beginning of the proof. Therefore, the size of $|P'|$ can be bounded by $|P'| \leq 3|P|^2$, since the size of the proofs from Lemma 10 are strictly less than $3|P|$.

The proof P' has all cut formulas in Σ_d or Π_d, where $d = \mathsf{aqd}(P)$. It suffices to assume $d > 0$. Applying Theorem 5 d times gives a tree-like proof P'' with the same end cedent, in which all cut formulas are quantifier free, with $h(P'') \leq 2_{d-1}^{3|P|^2}$. Now, applying Theorem 8 of [5] and the discussion from the end of Section 4 of [5], we get a proof P''' of the same end cedent with height bounded by $h(P''') \leq 2^{|P|} 2_{d-1}^{3|P|^2}$, such that all cut formulas in P''' are atomic. Then, applying Lemma 7 of [5], we get another proof P'''' again with the same end cedent, which is cut free, and has height bounded by $2^{h(P''')+1}$. In particular, the size of P'''' is less than $2_2^{h(P''')+1}$.

Therefore, $|P''''| < 2_{d+2}^{|P|}$, at least for $|P| > 7$. For $d > 0$, this gives $|P''''| < 2_{d+2}^{|P|}$ for $|P| > 7$. This completes the proof of Theorem 9. $\square$

The size bound on P''' is not optimal; we expect that even $2_{d+1}^{|P|}$ might work.

Acknowledgements. We thank G. Mints for encouraging comments, and the referee for helpful comments and corrections.

References

[1] K. Aehlig and A. Beckmann, *On the computational complexity of cut reduction*, Annals of Pure and Applied Logic, 161 (2010), pp. 711–736.

[2] M. Baaz and A. Leitsch, *Cut normal forms and proof complexity*, Annals of Pure and Applied Logic, 97 (1999), pp. 127–177.

[3] A. Beckmann and S. R. Buss, *Separation results for the size of constant-depth propositional proofs*, Annals of Pure and Applied Logic, 136 (2005), pp. 30–55.

[4] ——, *Characterization of definable search problems in bounded arithmetic via proof notations*, in Ways of Proof Theory, Ontos Verlag, 2010, pp. 65–134.

[5] ——, *Corrected upper bounds for free-cut elimination*, Theoretical Computer Science, 412 (2011), pp. 5433–5445.

[6] S. R. Buss, *An introduction to proof theory*, in Handbook of Proof Theory, S. R. Buss, ed., North-Holland, 1998, pp. 1–78.

[7] P. Gerhardy, *Refined complexity analysis of cut elimination*, in Proc. 17th Workshop on Computer Science Logic (CSL), Lecture Notes in Computer Science #2803, Springer Verlag, 2003, pp. 212–225.

[8] J.-Y. Girard, *Proof Theory and Logical Complexity*, Humantities Press, 1987.

[9] J. Krajíček, *Lower bounds to the size of constant-depth propositional proofs*, Journal of Symbolic Logic, 59 (1994), pp. 73–86.

[10] ——, *Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic*, Journal of Symbolic Logic, 62 (1997), pp. 457–486.

[11] G. Kreisel, G. E. Mints, and S. G. Simpson, *The use of abstract language in elementary metamathematics: Some pedagogical examples*, in Logic Colloquium, Lecture Notes in Mathematics #453, Springer-Verlag, 1975, pp. 38–131.

[12] G. E. Mints, *Finite investigations of transfinite derivations*, Journal of Soviet Mathematics, 10 (1978), pp. 548–596. Russian original in Zapiski Nauchnykh Seminarov LOMI Steklova AN SSSR 49 (1975) 67-121.

[13] V. P. Orevkov, *Upper bound on the lengthening of proofs by cut elimination*, Journal of Soviet Mathematics, 34 (1986), pp. 1810–1819. Original Russian version in Zap. Nauchn. Sem. L.O.M.I. Steklov 137 (1984) 87-98.

[14] ——, *Applications of cut elimination to obtain estimates of proof lengths*, Soviet Mathematics Doklady, 36 (1988), pp. 292–295. Original Russian version in Dokl. Akad. Nauk. 296, 3 (1987) 539-542.

[15] A. A. Razborov, *On provably disjoint NP-pairs*, Tech. Report RS-94-36, Basic Research in Computer Science Center, Aarhus, Denmark, November 1994. http://www.brics.dk/index.html.

[16] A. S. Troelstra and H. Schwichtenberg, *Basic Proof Theory*, Tracts in Theoretical Computer Science #43, Cambridge University Press, Cambridge, 2nd ed., 2000.

[17] W. ZHANG, *Cut elimination and automatic proof procedures*, Theoretical Computer Science, 91 (1991), pp. 265–284.

[18] ——, *Depth of proofs, depth of cut-formulas, and complexity of cut formulas*, Theoretical Computer Science, 129 (1994), pp. 193–206.