On Transformations of Constant Depth Propositional Proofs

Arnold Beckmann[a], Sam Buss[b,1]

[a] *Department of Computer Science, Swansea University, Swansea SA2 8PP, UK*
[b] *Department of Mathematics; University of California, San Diego; La Jolla, CA 92093-0112, USA*

Abstract

This paper studies the complexity of constant depth propositional proofs in the cedent and sequent calculus. We discuss the relationships between the size of tree-like proofs, the size of dag-like proofs, and the heights of proofs. The main result is to correct a proof construction in an earlier paper about transformations from proofs with polylogarithmic height and constantly many formulas per cedent.

Keywords: Lengths of proofs; Propositional calculus; Frege system
2010 MSC: primary 03F20; secondary 03F07, 68Q15, 68R99

---

1. Introduction

Constant depth Frege systems d-PK are propositional proof systems with a constant upper bound d on the alternation depth of unbounded fanin conjunctions and disjunctions. These systems have been extensively studied because of their connections to constant depth Boolean circuits and to first-order fragments S^i_2 and T^i_2 of bounded arithmetic; see [10, 1, 3, 6, 8, 9, 11, 4]. In [5], the present authors gave a synthesis and summary of constructions for constant depth PK proofs studying different notions of proof size under different assumptions about proof complexity. The paper [5] studied the size of proofs measured in terms of numbers of symbols and in terms of numbers of cedents (lines); it considered tree-like proofs, sequence-like (dag-like) proofs, proofs with constantly many formulas per cedent, and proofs of restricted height. That paper gave a comprehensive discussion of how proofs can be transformed with respect to these measures of proof complexity.

Unfortunately, one of the constructions in the previous paper was incorrect, namely Lemma 5 of [5] about converting tree-like depth $d+1$ PK proofs with constantly many formulas per cedent into tree-like depth d PK proofs with

[1] Supported in part by NSF grants CCF-121351 and DMS-1101228, and a Simons Foundation Fellowship 306202.

only a polynomial blowup in proof size. This is important (and is based on ideas from Razborov [12] and Krajíček [7]), because it can be a crucial step in transforming proofs obtained by the Paris-Wilkie translation from first-order bounded arithmetic proofs into lower depth Frege proofs. The main result of the present paper is to give a corrected construction, establishing nearly all of the results claimed in [5].

The results in [5] were stated first for quasipolynomial size proofs and then for polynomial size proofs. The quasipolynomial bounds are usually the most relevant for applications to bounded arithmetic, since this is what comes from the Paris-Wilkie translations of fragments of T_2. They are also usually the most relevant way to measure the size of constant depth Boolean circuits, see [2], since quasipolynomial size arises naturally from the Furst-Saxe-Sipser translations of polynomial time hierarchy predicates. Our corrected constructions establish all of the results claimed for quasipolynomial size PK proofs in [5].

The second set of results claimed in [5] concerned polynomial size proofs. We are able to give corrected proofs of part of these results, namely the parts relating sequence size and tree size for propositional proofs. However, we are unable to prove the part of these results which concerned height restricted proofs. As a replacement, we state (in Corollary 12) a connection to a different restriction of propositional proofs which is useful for obtaining polynomial size propositional proofs from, say, fragments of the bounded arithmetic theories $\mathrm{I}\Delta_0$ or T_1.

The remainder of the introduction states the key definitions about constant depth proofs and (re)introduces notations. The reader can refer to [5] for further discussions and definitions. After that, we state the main results claimed in [5]. Section 2 explains the error in the earlier construction, and how to fix it in the setting of quasipolynomial size proofs. Section 3 then gives the necessary additional arguments to establish the results about polynomial size Frege proofs.

We thank an anonymous referee of one of our other papers for pointing out the error in [5].

1.1. Definitions of proof systems

We follow the notations from [5], with the sole exception that we use "PK" instead of "LK" to emphasize the fact that the proof systems are propositional. We also make some minor inessential changes to the way formulas and derivations are defined. We work in classical propositional logic, essentially the sequent calculus, but formulated for simplicity using Tait-style proofs in which the lines, called "cedents", are sets of formulas.

Propositional formulas are built from propositional variables p_i, negated variables $\neg p_i$, and unbounded fanin conjunctions $\bigwedge$ and disjunctions $\bigvee$. All formulas are classified as being "literals" or "$\bigvee$-formulas" or "$\bigwedge$-formulas". The *literals* are the variables and negated variables p_i and $\neg p_i$. The $\bigvee$- and $\bigwedge$-formulas are defined inductively by

- If Φ is a nonempty finite set of literals and $\bigvee$-formulas, then $\bigwedge\Phi$ is an $\bigwedge$-formula.

- If Φ is a nonempty finite set of literals and $\bigwedge$-formulas, then $\bigvee\Phi$ is an $\bigvee$-formula.

The point of these definitions is that adjacent $\bigwedge$'s (respectively, $\bigvee$'s) must be collapsed; for instance, an $\bigwedge$-formula cannot be an argument to an $\bigwedge$.

We define $\wedge$ to be a binary operation on formulas. If φ and ψ are *not* $\bigwedge$-formulas, then $\varphi \wedge \psi$ is the formula $\bigwedge\{\varphi, \psi\}$; in addition, $\varphi \wedge \bigwedge\Psi$ denotes $\bigwedge(\{\varphi\} \cup \Psi)$, and $\bigwedge\Phi \wedge \bigwedge\Psi$ denotes $\bigwedge(\Phi \cup \Psi)$. The binary operation $\vee$ is defined similarly.

For φ a formula, $\neg\varphi$ abbreviates the formula formed from φ by interchanging $\bigwedge$ and $\bigvee$, and interchanging atoms and their negations.

For an arbitrary nonempty finite set of formulas Φ, which may contain $\bigvee$-formulas, let $\bigvee^{\mathrm{coll}}\Phi$ abbreviate the formula which is built from $\bigvee\Phi$ by collapsing any adjacent $\bigvee$'s. Formally, let Φ^{coll} be the set of formulas φ such that φ is in Φ but not an $\bigvee$-formula, or $\varphi \in \Psi$ for some formula $\bigvee\Psi$ in Φ. Define $\bigvee^{\mathrm{coll}}\Phi$ as $\bigvee\Phi^{\mathrm{coll}}$.

The *depth* $\mathrm{dp}(\varphi)$ of a formula φ is the maximal nesting of $\bigwedge$ and $\bigvee$ in φ. Thus, literals have depth 0, and $\mathrm{dp}(\bigwedge\Phi) = 1 + \max\{\mathrm{dp}(\varphi) : \varphi \in \Phi\}$, etc.

A line in a PK-proof is a finite set of formulas called a *cedent*. We use capital Greek letters $\Gamma, \Delta, \ldots$ to denote cedents. The intended meaning of a cedent Γ is $\bigvee\Gamma$. Cedents are sometimes also called *clauses* (in the case of refutations). We often abuse notation by writing Γ, φ or $\Gamma \vee \varphi$ instead of $\Gamma \cup \{\varphi\}$, or by writing $\varphi_1, \ldots, \varphi_k$ instead of $\{\varphi_1, \ldots, \varphi_k\}$, etc.

A set $\mathcal{A}$ of *nonlogical axioms* is a set of cedents. The intended meaning of $\mathcal{A}$ is the conjunction its members. The axioms and rules of inference for PK are as follows. We write $\mathrm{card}(X)$ for the cardinality of a set X.

Logical axioms: Any cedent $\varphi, \neg\varphi$ for φ a literal.

Nonlogical axioms: The cedent Γ and the cedent $\bigvee^{\mathrm{coll}}\Gamma$ for any $\Gamma \in \mathcal{A}$.

$\bigvee$-introduction: For $\varphi \in \Phi$, and $\bigvee\Phi$ a formula: $\quad \bigvee \dfrac{\Gamma, \varphi}{\Gamma, \bigvee\Phi}$

$\bigwedge$-introduction: For $\bigwedge\Phi$ a formula, the inference with $\mathrm{card}(\Phi)$ many premises:

$$\bigwedge \frac{\Gamma, \varphi \qquad \text{for } \varphi \in \Phi}{\Gamma, \bigwedge\Phi}$$

Structural rules: The weakening rule and the cut rule:

$$\text{weakening} \ \frac{\Gamma}{\Gamma, \Gamma'} \qquad\qquad \text{cut} \ \frac{\Gamma, \neg\varphi \qquad \Gamma, \varphi}{\Gamma}$$

The formula φ in the premises of the cut-rule is called the *cut-formula* of this rule. Following [5], we allow both Γ and $\bigvee^{\mathrm{coll}}\Gamma$ as nonlogical axioms; this is reasonable enough as they have the same meaning. This convention does not affect proof size substantially, and it simplified some technical aspects in our constructions.

A PK-*derivation from* $\mathcal{A}$ is a tree in which each node is labeled with a cedent. We picture the tree with the root at the bottom. Cedents at leaf nodes must be logical axioms or nonlogical axioms from $\mathcal{A}$, and cedents at internal nodes are inferred by one of the rules of inference from the cedents on the children. If the root is labeled with Γ, the derivation is a *derivation of* Γ *from* $\mathcal{A}$. If the root is labelled with the empty cedent, then it is a PK *refutation of* $\mathcal{A}$.

The complexity of a derivation π can be measured in several ways. We use just "size" when counting the number of occurrences of symbols in the derivation, and "cedent size" to count the number of occurrences of cedents. We use the adjectives "tree" and "sequence" to denote whether the proof is tree-like, or is to be converted to a dag without repetition of cedents. The *tree-cedent-size* of π is the number of cedents (hence, nodes) in the tree π; the *sequence-cedent-size* is the number of distinct cedents in π. The *tree-size* is the total number of symbol occurrences in cedents in π; the *sequence-size* is the total number of symbol occurrences in distinct cedents in π. The *height* of π is the maximum number of cedents along any path in π.

Constant depth PK proof systems are defined by restricting the depth of formulas appearing in refutations. We use a parameter $S > 0$ to bound both the fanin of $\bigwedge$'s and $\bigvee$'s, and the size of derivations. It is useful to treat $\bigwedge_{i<\log S}\varphi_i$ and $\bigvee_{i<\log S}\varphi_i$ for literals φ_i as being depth 1/2. This motivates the following definition, which generalizes both the usual notion of depth and Krajíček's notion of "Σ-depth" [6].

Definition 1. Let $S \in \mathbb{N}$. The classes Θ_d^S for $d \in \frac{1}{2}\mathbb{N} = \{0, \frac{1}{2}, 1, 1\frac{1}{2}, 2, \ldots\}$ are defined inductively as follows:

1. $\varphi \in \Theta_0^S$ iff φ is a literal.
2. $\varphi \in \Theta_{0.5}^S$ iff φ is a literal or an $\bigwedge$ or $\bigvee$ of at most $\log S$ many literals.
3. $\varphi \in \Theta_{d+1}^S$ iff φ is in Θ_d^S, or it is an $\bigvee$ or $\bigwedge$ of at most S many formulas from Θ_d^S.

Observe that for all $d \in \frac{1}{2}\mathbb{N}$, Θ_d^S is a subset of $\Theta_{d+\frac{1}{2}}^S$, which can be seen by induction on d. We say that $\varphi \in \Theta_d$ if $\varphi \in \Theta_d^S$ for some S. A cedent Γ is in Θ_d^S or Θ_d iff all the formulas in Γ are in Θ_d^S or Θ_d, respectively. A derivation is in Θ_d^S or Θ_d iff it contains only cedents in Θ_d^S or Θ_d, respectively.

We refer to $\varphi \in \Theta_d$ as φ *being of depth* d.

Definition 2. Let "X-size" mean one of our four size measurements "tree-size", "sequence-size", "tree-cedent-size", or "sequence-cedent-size". A derivation π is a d-PK *derivation of X-size* S if it has X-size $\leq S$ and all formulas appearing in π are in Θ_d^S.

Our main theorems are asymptotic results about refutations of families $\{\mathcal{A}_n\}_n$ of formulas. A size bound S is *polynomial*, respectively *quasipolynomial*, if it is of the form $S = n^{O(1)}$, respectively, $S = 2^{(\log n)^{O(1)}}$.

1.2. Main Theorems

Our theorems concern (quasi)polynomial size derivations which are (a) dag-like, or (b) tree-like, or (c) tree-like height logarithmically bounded in terms of size. The results state that going from (a) to (b), or from (b) to (c) can be done at a cost of increasing formula depth by 1. Conversely, formula depth can be decreased by 1 by going from (c) to (b), or from (b) to (a).

The first main result claimed in [5] was:

Theorem 3. (Theorem 2 in [5]) *Let* $d \in \frac{1}{2}\mathbb{N}$*. For* $n \in \mathbb{N}$*, let* $\mathcal{A}_n$ *be a collection of* Θ_d*-cedents. Then the following conditions* (1) *and* (2) *are equivalent:*

(1) $\mathcal{A}_n$ *has a* d-PK *refutation of sequence-size quasipolynomial in* n*, for all* n*.*

(2) $\mathcal{A}_n$ *has a* $(d{+}1)$-PK *refutation of tree-size quasipolynomial in* n*, for all* n*.*

Furthermore, the following conditions (3) *and* (4) *are equivalent:*

(3) $\mathcal{A}_n$ *has a* d-PK *refutation of tree-size quasipolynomial in* n*, for all* n*.*

(4) $\mathcal{A}_n$ *has a* $(d{+}1)$-PK *refutation which simultaneously has tree-size quasipolynomial in* n *and height polylogarithmic in* n*, for all* n*.*

Note that (3) is the same as (2) except using d-PK instead of $(d{+}1)$-PK. Thus (1) and (2) are also equivalent to

(4′) $\mathcal{A}_n$ has a $(d{+}2)$-PK refutation which simultaneously has tree-size quasipolynomial in n and height polylogarithmic in n, for all n.

The error in [5] affects the proof of the implications from (4) to (3) and from (4′) to (2). This is corrected in Section 2.

The analogue of Theorem 3 for polynomial size growth rates was also stated in [5]. Here is the version which we are presently able to prove:

Theorem 4. (Adapted from Theorem 10 of [5]) *Let* $d \in \frac{1}{2}\mathbb{N}$*,* $n \in \mathbb{N}$*, and* $\mathcal{A}_n$ *a collection of* Θ_d*-cedents. Then the following conditions* (1) *and* (2) *are equivalent:*

(1) $\mathcal{A}_n$ *has a* d-PK *refutation of sequence-size polynomial in* n*, for all* n*.*

(2) $\mathcal{A}_n$ *has a* $(d{+}1)$-PK *refutation of tree-size polynomial in* n*, for all* n*.*

Theorem 4 is proved in Section 3.

The difference between the above Theorem 4 and Theorem 10 of [5] is that the latter also claimed an equivalence between d-PK refutations of polynomial tree-size and $(d{+}1)$-PK refutations of simultaneous polynomial tree-size and logarithmic height with cedents containing constantly many formulas:

Former Statement 5. (From Theorem 10 in [5]) *For* $d \in \frac{1}{2}\mathbb{N}$*,* $n \in \mathbb{N}$*, and* $\mathcal{A}_n$ *a collection of* Θ_d*-cedents, the following conditions* (3) *and* (4) *are equivalent:*

(3) $\mathcal{A}_n$ *has a* d-PK *refutation of tree-size polynomial in* n*, for all* n*.*

(4) $\mathcal{A}_n$ *has a* $(d{+}1)$-PK *refutation which simultaneously has tree-size polynomial in* n*, has height logarithmic in* n*, and has* $O(1)$ *many formulas in each cedent, for all* n*.*

We are unable to prove this and leave it here as an open problem. Corollary 12 provides a partial replacement.

2. Quasipolynomial size proofs

We are also unable to correct the proof of the next statement. It will be replaced by Theorem 8 and Corollary 10 below.

Former Statement 6. (Lemma 5 in [5]) *Let* $d \in \frac{1}{2}\mathbb{N}$*. Assume* $\mathcal{A}$ *is a collection of* Θ_d^S*-cedents, and* $\mathcal{A}$ *has a* Θ_{d+1}^S-PK *refutation* R *of tree-cedent-size* $\le S$*, where each cedent in the refutation consists of at most* λ *many formulas. Then* $\mathcal{A}$ *has a* Θ_d^S-PK *refutation of tree-cedent-size* $\le S^{\lambda+1}$*.*

The erroneous proof in [5] was based on the following construction, which will also be used for our corrected proof of Theorem 3. Any cedent Γ in the PK refutation R can be uniquely written in the form $\Delta \cup \Sigma \cup \Pi$ where Δ is $\Gamma \cap \Theta_d^S$, and where Σ and Π contain only $\bigvee$-formulas and $\bigwedge$-formulas (respectively) of depth $> d$. The subcedent Σ contains N formulas and equals $\{\bigvee_{j<n_i} A_{i,j} : i < N\}$ for integers $n_i > 0$, where the $A_{i,j}$'s are literals or $\bigwedge$-formulas of depth $\le d$. The subcedent Π contains M formulas and is equal to $\{\bigwedge_{j<m_i} B_{i,j} : i < M\}$ for integers $m_i > 0$, where the $B_{i,j}$'s are literals or $\bigvee$-formulas of depth $\le d$. Each cedent Γ is replaced by the collection of cedents Γ_f defined as

$$\Delta \cup \{A_{i,j} : i < N, j < n_i\} \cup \{B_{i,f(i)} : i < M\}, \tag{1}$$

for all functions f such that $f(i) < m_i$ for all $i < M$. The cedents Γ_f are evidently all in Θ_d^S, and there are $m_1 m_2 \cdots m_M \le S^\lambda$ many of them, because $M \le \lambda$ and since each $m_i \le S$.

It was claimed in [5] that the cedents Γ_f could be combined in a tree-like way so as to give a new depth d refutation. However, this claim is not true in general. The straightforward way of combining the translated cedents requires some cedents to be used multiple times: this would be fine if we were constructing sequence-like proofs, but not for constructing tree-like proofs. To illustrate this, consider the following counterexample (suggested to us by a referee of another of our papers). Let $\varphi_i = A_{i,1} \wedge A_{i,2}$, so $\neg\varphi_i$ is $\neg A_{i,1} \vee \neg A_{i,2}$, and assume that π_i is a derivation of $\neg\varphi_i, \varphi_{i+1}$. Consider the following (tree-like) derivation of $\neg\varphi_0, \varphi_n$:

$$\begin{gathered} \dfrac{\dfrac{\overset{\pi_1}{\neg\varphi_0,\varphi_1} \qquad \overset{\pi_2}{\neg\varphi_1,\varphi_2}}{\neg\varphi_0,\varphi_2} \qquad \overset{\pi_3}{\neg\varphi_2,\varphi_3}}{\neg\varphi_0,\varphi_3} \\ \ddots \\ \dfrac{\neg\varphi_0,\varphi_{n-1} \qquad \overset{\pi_n}{\neg\varphi_{n-1},\varphi_n}}{\neg\varphi_0,\varphi_n} \end{gathered} \tag{2}$$

The above construction of the Γ_f's replaces each of the topmost subderivations π_i, for $i = 1, \ldots, n$, with two derivations π_i^j of $\neg A_{i-1,1}, \neg A_{i-1,2}, A_{i,j}$, for $j = 1, 2$. These are combined to form derivations $\widehat{\pi}_i^j$ of $\neg A_{0,1}, \neg A_{0,2}, A_{i,j}$ as follows. First, $\widehat{\pi}_1^j$ is π_1^j. Second, inductively define $\widehat{\pi}_{i+1}^j$ by replacing the cut inference

$$\frac{\neg\varphi_0, \varphi_i \qquad \neg\varphi_i, \varphi_{i+1}}{\neg\varphi_0, \varphi_{i+1}} \tag{3}$$

of the derivation (2) with two cut inferences to form $\widehat{\pi}_{i+1}^j$:

$$\dfrac{\begin{matrix}\widehat{\pi}_i^1 \\ \neg A_{0,1}, \neg A_{0,2}, A_{i,1}\end{matrix} \qquad \dfrac{\begin{matrix}\widehat{\pi}_i^2 \\ \neg A_{0,1}, \neg A_{0,2}, A_{i,2}\end{matrix} \qquad \begin{matrix}\pi_{i+1}^j \\ \neg A_{i,1}, \neg A_{i,2}, A_{i+1,j}\end{matrix}}{\neg A_{0,1}, \neg A_{0,2}, \neg A_{i,1}, A_{i+1,j}}}{\neg A_{0,1}, \neg A_{0,2}, A_{i+1,j}} \tag{4}$$

The problem is that the derivations $\widehat{\pi}_i^1$ and $\widehat{\pi}_i^2$ are used twice: once to form $\widehat{\pi}_{i+1}^1$ and once to form $\widehat{\pi}_{i+1}^2$. The result is that at each inductive step, the numbers of occurrences π_1^1 and π_1^2 get doubled (and similarly for other subderivations inside $\widehat{\pi}_i^1$ and $\widehat{\pi}_i^2$). Thus, the final tree-cedent-sizes of the derivations $\widehat{\pi}_n^j$ are exponential in n instead of polynomial in n.

Although this is a counterexample to the proof idea, it is not a counterexample to the Former Statement 6. Indeed, the cuts in the derivation (2) can be rearranged in a way such that the blowup as experienced above does not occur:

$$\dfrac{\begin{matrix}\pi_1 \\ \neg\varphi_0, \varphi_1\end{matrix} \qquad \dfrac{\begin{matrix}\pi_2 \\ \neg\varphi_1, \varphi_2\end{matrix} \qquad \begin{matrix} \dfrac{\begin{matrix}\pi_{n-1} \\ \neg\varphi_{n-2}, \varphi_{n-1}\end{matrix} \qquad \begin{matrix}\pi_n \\ \neg\varphi_{n-1}, \varphi_n\end{matrix}}{\neg\varphi_{n-2}, \varphi_n} \\ \iddots \\ \neg\varphi_2, \varphi_n\end{matrix}}{\neg\varphi_1, \varphi_n}}{\neg\varphi_0, \varphi_n}$$

It is not hard to check that in this case the above construction does not cause an exponential blowup in size; in fact, a polynomial size tree-like proof is obtained.

The point is that the unwanted doubling in the above example occurred asymmetrically. The inference (3) was twice transformed into inferences (4), once for $j = 1$ and once for $j = 2$. This caused a double use of the derivations $\widehat{\pi}_i^1$ and $\widehat{\pi}_i^2$, but did not cause a double use of the derivations π_{i+1}^1 and π_{i+1}^2. This motivates the following definitions.

Fix a value for d, and let π be a Θ_{d+1}^S-derivation. Suppose

$$\frac{\Gamma, C \qquad \Gamma, \neg C}{\Gamma}$$

is a cut inference in which C is a depth $> d$ $\bigwedge$-formula, and thus $\neg C$ is a depth $> d$ $\bigvee$-formula. Then Γ, C is called a *critical* cedent. (It does not matter whether Γ, C is written as the left or right premiss.) Any depth $> d$ $\bigwedge$-formula (such as C) is called a *critical* formula.

Definition 7. The *critical cut rank* of a derivation π is equal to the maximum number of critical cedents on any path from the root of π to any axiom in π.

The intuition is that the critical cut rank of a cedent controls the amount of doubling that occurs when applying the above transformation for the Former Statement 6. This is formalized by the next theorem.

Theorem 8. *Let $\mathcal{A}$ be a collection of Θ_d^S-cedents. Suppose $\mathcal{A}$ has a Θ_{d+1}^S-*PK *refutation π of tree-size $\le S$ and critical cut rank ρ. Then $\mathcal{A}$ has a Θ_d^S-*PK *refutation of tree-size $\le O(S^{\rho+3})$.*

There is no need for Theorem 8 to assume any upper bound λ on the number of formulas in each cedent. This is because the critical cut rank provides an implicit bound on the number of critical formulas in each cedent. That is, the number M of $\bigwedge$-formulas of depth $> d$ is bounded by ρ, as each critical formula must be eliminated by a critical cut.

Theorem 8 and Corollary 10 below are stated in terms of "tree-size"; the same proof method also gives similar statements for "tree-cedent-size".

Lemma 9. *Let A be a formula of size S.*

a. *The cedent $A, \neg A$ has a cut-free* PK*-derivation of, simultaneously, tree-cedent-size $O(S)$ and tree-size $O(S^2)$.*

b. *Assume $A = \bigvee\Phi$. The cedent $\neg A, \Phi$ has a cut-free* PK*-derivation of, simultaneously, tree-cedent-size $O(S)$ and tree-size $O(S^2)$.*

Proof of Lemma 9. The proof is entirely standard, so we mostly omit it. One first proves part a. by induction on the size of A. Part b. then follows easily. We prove only part b. as an example. For each $\varphi \in \Phi$, part a. gives a PK-derivation of $\neg\varphi, \varphi$; then a weakening gives $\neg\varphi, \Phi$. The formula $\neg A$ is the conjunction of the $\neg\varphi$'s for $\varphi \in \Phi$, so one more $\bigwedge$-inference gives $\neg A, \Phi$ as desired. □

Proof of Theorem 8. We use exactly the same construction as was described for the Former Statement 6. Each cedent Γ in the derivation is expressed exactly as before in the form Δ, Σ, Π; the $A_{i,j}$'s, the $B_{i,j}$'s, N, M, the n_i's and the m_i's are exactly as before. We again let Γ_f denote the cedent (1).

Let π_Γ be the subderivation of π ending with Γ, and $\mathrm{tcs}(\pi_\Gamma)$ be the tree-cedent-size of π_Γ. Let $\mathcal{A}^{\mathrm{coll}} = \{\Phi^{\mathrm{coll}} : \Phi \in \mathcal{A}\}$ be the set of collapsed axioms, where disjunctions in cedents are replaced by the set of disjuncts.

Claim. *Let ρ be the critical cut rank of the subderivation π_Γ. For f any function with $f(i) < m_i$ for all $i < M$, the cedent Γ_f has a Θ_d^S-derivation of tree-cedent-size $\le 2 \cdot S^\rho \cdot \mathrm{tcs}(\pi_\Gamma)$ from the nonlogical axioms $\mathcal{A}$ and $\mathcal{A}^{coll}$. Each cedent in this derivation has size $O(S)$.*

The claim implies Theorem 8, since the final line of the refutation π is the empty cedent $\emptyset$, and its translation $\emptyset_f$ is just the empty cedent. Each axiom in $\mathcal{A}^{\mathrm{coll}}$ has a Θ_d^S-derivation of tree-cedent-size $\le S$, in the following way:

Let $\Phi \in \mathcal{A}$ be written in the form $\Psi_0, \bigvee\Psi_1, \ldots, \bigvee\Psi_k$ with the Ψ_i's consisting of literals and $\bigwedge$-formulas. Thus $\Phi^{\rm coll}$ is the cedent $\Psi_0, \Psi_1, \ldots, \Psi_k$. Using Lemma 9 we can derive $\neg\bigvee\Psi_i, \Psi_i$ for all $i > 0$. Using these and weakening and the axiom Φ we obtain with successive cuts $\Psi_0, \bigvee\Psi_1, \ldots, \bigvee\Psi_i, \Psi_{i+1}, \ldots, \Psi_k$ for $i = k, \ldots, 0$, deriving $\Phi^{\rm coll}$ when $i = 0$. Adding these derivations of axioms in $\mathcal{A}^{\rm coll}$ requires only $O(S)$ steps per axiom; in addition, $\mathrm{tcs}(\pi_\Gamma) < S$. Thus, the refutation of $\mathcal{A}$ has $O(S^{\rho+2})$ occurrences of cedents, each of size $O(S)$. Theorem 8 follows.

The proof of the claim proceeds by induction on the number of cedents in π_Γ. The base case of the induction is when Γ is either a logical or nonlogical axiom. In this case, the critical cut rank ρ equals 0, and either the cedent Δ is all of Γ and both Σ and Π are empty, or Σ is $\bigvee^{\rm coll}\Phi$ for some Φ in $\mathcal{A}$ and Δ and Π are empty. In either case, Π is empty, so there is only one cedent Γ_f — with f the empty function — and it is equal to Γ in the former case, or $\Phi^{\rm coll}$ in the latter.

The induction cases where Γ is inferred from a cedent Γ' by either a weakening or an $\bigvee$-introduction are trivial: the critical cut rank of π_Γ is equal to the critical cut rank of $\pi_{\Gamma'}$ of course. For any appropriate function f, either $\Gamma_f = \Gamma'_f$ or Γ_f can be derived from Γ'_f by a weakening or $\bigvee$-introduction.

The induction case for an $\bigwedge$-introduction rule which introduces a formula of depth $\leq d$ is handled just like the previous cases. Suppose Γ has the form $\Gamma_0, \bigwedge\{B_{M,j} : j < m_M\}$ and is inferred by the m_M premiss inference

$$\frac{\Gamma_0, B_{M,j} \qquad\qquad j < m_M}{\Gamma_0, \bigwedge\{B_{M,j} : j < m_M\}}$$

where the $\bigwedge$-formula has depth $> d$. There are M many depth $> d$ $\bigwedge$-formulas in Γ_0. A function f suitable for defining Γ_f has domain $[M{+}1] = \{0, \ldots, M\}$. Let $\Gamma_{f(M)}$ be the $f(M)$-th premiss of the $\bigwedge$-introduction; namely, $\Gamma_{f(M)}$ is $\Gamma_0, B_{M,f(M)}$. The translation Γ_f of Γ is the same cedent as the translation $(\Gamma_{f(M)})_{f'}$ of $\Gamma_{f(M)}$, where $f' = f{\restriction}[M]$ is the restriction of f to the domain $[M] = \{0, \ldots, M{-}1\}$. The critical cut rank does not increase when going to a premiss, so the induction hypothesis gives a derivation of Γ_f of tree-cedent-size $\leq 2 \cdot S^\rho \cdot \mathrm{tcs}(\pi_{\Gamma_{f(M)}}) < 2 \cdot S^\rho \cdot \mathrm{tcs}(\pi_\Gamma)$.

The case of a cut on a formula of depth d is also handled similarly to the cases of weakening and $\bigvee$-introduction. Finally suppose Γ is inferred by a cut on a depth $> d$ formula:

$$\frac{\Gamma, \bigwedge\{B_{M,j} : j < m_M\} \qquad \Gamma, \bigvee\{\neg B_{M,j} : j < m_M\}}{\Gamma}$$

Let π_1 and π_2 be the subderivations of the left and right premisses, respectively, and let S_1 and S_2 be their tree-cedent-sizes. We have $\mathrm{tcs}(\pi_\Gamma) = S_1 + S_2 + 1$. The critical cut rank of π_1 is $\leq \rho{-}1$, and the critical cut rank of π_2 is $\leq \rho$ (and equality holds in at least one of these cases). Fix any suitable function f with domain $[M]$ for defining a translation Γ_f of Γ. The induction hypothesis for the right premiss gives a Θ^S_d-derivation of

$$\Gamma_f, \neg B_{M,0}, \neg B_{M,1}, \ldots, \neg B_{M,m_M-1} \tag{5}$$

of tree-cedent-size $\leq 2 \cdot S^\rho \cdot S_2$. Now, for each $i < m_M$, let f_i be the function which extends f to the domain $[M{+}1]$ with the value $f_i(M) = i$. The induction hypothesis applied to the left premiss using the function f_i gives a Θ_d^S-derivation of

$$\Gamma_f, B_{M,i} \tag{6}$$

of tree-cedent-size $\leq 2 \cdot S^{\rho-1} \cdot S_1$. Combining the cedent (5) and the m_M cedents (6) for $i < m_M$ using m_M cut inferences gives a Θ_d^S-derivation of Γ_f of tree-cedent-size bounded by (since $m_M < S$)

$$\begin{aligned}
&2 \cdot (S^\rho \cdot S_2) + 2 \cdot m_M \cdot (S^{\rho-1} \cdot S_1) + 2 \cdot m_M \\
&\quad \leq \quad 2 \cdot (S^\rho \cdot S_2) + 2 \cdot S \cdot (S^{\rho-1} \cdot S_1) + 2 \cdot S \\
&\quad \leq \quad 2 \cdot S^\rho (S_1 + S_2 + 1) \\
&\quad = \quad 2 \cdot S^\rho \cdot \mathrm{tcs}(\pi_\Gamma).
\end{aligned}$$

Note that the factor of 2 in "$2 \cdot m_M$" comes from the fact that the cuts must be proceeded by weakening inferences to make the side formulas match. This completes the proof of the claim, and thereby Theorem 8. □

Since the critical cut rank of a derivation is bounded by its height, Theorem 8 immediately implies:

Corollary 10. *Assume $\mathcal{A}$ is a collection of Θ_d^S-cedents and $\mathcal{A}$ has a Θ_{d+1}^S-*PK *refutation of tree-size $\leq S$ and height h. Then $\mathcal{A}$ has a Θ_d^S-*PK *refutation of tree-size $\leq S^{h+3}$.*

With this corollary, we are now ready to close the gap in the proof of Theorem 3. We give a sketch of the complete proof, relying on statements from [5], but using Corollary 10 instead of Lemma 5 in [5] (the Former Statement 6). We first repeat the relevant statements from [5].

Lemma (Lemma 4 in [5])**.** *Assume $\mathcal{A}$ has a Θ_d^S-*PK *refutation R of sequence-size $\leq S$. Then $\mathcal{A}$ has a Θ_{d+2}^S-*PK *refutation R' which is simultaneously of height $\log S + O(1)$ and tree-size $O(S^4)$. Furthermore, each cedent in R' has $O(1)$ many formulas.*

Lemma (Lemma 6 in [5])**.** *Let $\mathcal{A}$ be a collection of Θ_d^S-cedents. If $\mathcal{A}$ has a Θ_{d+1}^S-*PK *refutation of tree-size S, then $\mathcal{A}$ has a Θ_d^S-*PK *refutation of sequence-size $3S^2$.*

Corollary (Corollary 9 in [5])**.** *Let $\mathcal{A}$ be a collection of Θ_d^S-cedents. Suppose $\mathcal{A}$ has a Θ_d^S-*PK *refutation of tree-size S. Then $\mathcal{A}$ has a Θ_{d+1}^S-*PK *refutation of height $O(\log S)$.*

Proof of Theorem 3. To prove (1) and (2) are equivalent, use

[5, Lemma 4] shows	(1) for d-PK	$\Longrightarrow$	(4$'$) for (d+2)-PK
[5, Lemma 6] shows	(2) for (d+1)-PK	$\Longrightarrow$	(1) for d-PK
Corollary 10 shows	(4$'$) for (d+2)-PK	$\Longrightarrow$	(2) for (d+1)-PK

To show that (3) and (4) are equivalent, use:

Corollary 10 shows	(4) for $(d{+}1)$-PK	$\Longrightarrow$	(3) for d-PK
[5, Corollary 9] shows	(3) for d-PK	$\Longrightarrow$	(4) for $(d{+}1)$-PK. $\square$

3. Polynomial size proofs

This section proves Theorem 4, and states and proves Corollary 12 about polynomial size simulations.

Proposition 11. *Let $\mathcal{A}$ be a collection of Θ_d^S-cedents. Suppose $\mathcal{A}$ has a Θ_d^S-*PK *refutation π of sequence-size $\leq S$. Then $\mathcal{A}$ has a Θ_{d+1}^S-*PK *refutation of tree-size $O(S^4)$.*

A similar statement holds in terms of cedent-size. Namely, a Θ_d^S-PK refutation of sequence-cedent-size $\leq S$ can be transformed into a Θ_{d+1}^S-PK refutation of tree-cedent-size $O(S^3)$.

Proof. We only give a sketch of the proof, which is based on a construction of Krajíček [6, Prop 1.1]. A refutation of $\mathcal{A}$ of sequence-size $\leq S$ can be written as a sequence of cedents $\Gamma_1, \ldots, \Gamma_L{=}\emptyset$ where each Γ_i is a logical axiom, is a nonlogical axiom Γ or $\bigvee^{\mathrm{coll}}\Gamma$ with $\Gamma \in \mathcal{A}$, or is formed from previous cedents $\Gamma_1, \ldots, \Gamma_{i-1}$ by applying one of the rules of PK. Let A_i be $\bigvee^{\mathrm{coll}}\Gamma_i$. We claim that each cedent $\neg A_1, \ldots, \neg A_{i-1}, A_i$ has a cut-free PK derivation from $\mathcal{A}$ of tree-size $O(S^3)$. Since the derivations are cut-free, they are Θ_d^S-derivations.

We will only consider the case that the last inference has been an $\bigwedge$-inference, as this is the hardest case. All other cases are similar and left to the reader.

Suppose that Γ_i is the cedent $\Delta, \bigwedge_{j<J}\varphi_j$ which is derived by an $\bigwedge$-inference from the cedents $\Gamma_{i_j} = \Delta, \varphi_j$. We show how to obtain the cut-free PK derivation of $\neg A_1, \ldots, \neg A_{i-1}, A_i$ from $\mathcal{A}$ of tree-size $O(S^3)$.

We have $A_i = \bigvee^{\mathrm{coll}}\Gamma_i = \bigvee\Gamma_i^{\mathrm{coll}}$ and $\Gamma_i^{\mathrm{coll}} = \Delta^{\mathrm{coll}} \cup \{\bigwedge_{j<J}\varphi_j\}$.

The condition for $\bigwedge_{j<J}\varphi_j$ being a formula implies that the φ_j's are either literals or $\bigvee$-formulas. W.l.o.g., we assume that each φ_j for $j < J_0$ is a literal, and that each φ_j for $J_0 \leq j < J$ is a $\bigvee$-formula of the form $\varphi_j = \bigvee_{k<K_j}\psi_{k,j}$. Then A_{i_j} is of the form $\bigvee^{\mathrm{coll}}\Gamma_{i_j} = \bigvee\Gamma_{i_j}^{\mathrm{coll}}$ with $\Gamma_{i_j}^{\mathrm{coll}} = \Delta^{\mathrm{coll}} \cup \{\varphi_j\}$ for $j < J_0$, and $\Gamma_{i_j}^{\mathrm{coll}} = \Delta^{\mathrm{coll}} \cup \{\psi_{k,j} \colon k < K_j\}$ for $J_0 \leq j < J$.

For each formula $\delta \in \Delta^{\mathrm{coll}}$ we form the following derivation $(\star_\delta)$:

$$\text{weakening}\,\dfrac{\bigvee\,\dfrac{\neg\delta, \delta}{\neg\delta, \bigvee^{\mathrm{coll}}\Gamma_i}}{\neg\delta, A_i, \varphi_j} \tag{$\star_\delta$}$$

The first inference uses $\delta \in \Delta^{\mathrm{coll}} \subseteq \Gamma_i^{\mathrm{coll}}$. The second inference uses the fact that A_i is the same as $\bigvee\Gamma_i^{\mathrm{coll}}$, and introduces φ_j by weakening.

For $j < J_0$, we form

$$\text{weakening}\ \frac{\neg\varphi_j, \varphi_j}{\neg\varphi_j, A_i, \varphi_j}$$

A single $\bigwedge$-introduction inference from these derivations and the $(\star_\delta)$'s for $\delta \in \Delta^{\mathrm{coll}}$ yields $\neg A_{i_j}, A_i, \varphi_j$, since

$$\neg A_{i_j} \;=\; \bigwedge\{\neg\gamma\colon \gamma \in \Gamma^{\mathrm{coll}}_{i_j}\} \;=\; \bigwedge\{\neg\delta\colon \delta \in \Delta^{\mathrm{coll}}\} \wedge \neg\varphi_j.$$

For $J_0 \le j < J$ and $k < K_j$, we form the following derivation $(\star_{j,k})$, using $\varphi_j = \bigvee_{k<K_j} \psi_{k,j}$:

$$\text{weakening}\ \frac{\bigvee\ \dfrac{\neg\psi_{k,j}, \psi_{k,j}}{\neg\psi_{k,j}, \varphi_j}}{\neg\psi_{k,j}, A_i, \varphi_j} \qquad (\star_{j,k})$$

A single $\bigwedge$-introduction inference from the $(\star_{j,k})$'s for $k < K_j$ and the $(\star_\delta)$'s for $\delta \in \Delta^{\mathrm{coll}}$ yields $\neg A_{i_j}, A_i, \varphi_j$, since

$$\neg A_{i_j} \;=\; \bigwedge\{\neg\gamma\colon \gamma \in \Gamma^{\mathrm{coll}}_{i_j}\} \;=\; \bigwedge\{\neg\delta\colon \delta \in \Delta^{\mathrm{coll}}\} \wedge \bigwedge\{\neg\psi_{k,j}\colon k < K_j\}.$$

The above gives, for all $j < J$, derivations of the cedents $\neg A_1, \ldots, \neg A_{i-1}, A_i, \varphi_j$ by using additional weakening inferences. By inspection, the tree-cedent-size of each derivation is order the size of A_{i_j}, that is $O(S)$. Also by inspection, the cedents in these derivations have size $O(S)$. Thus each of these derivations has tree-size $O(S^2)$.

The final step is to apply an $\bigwedge$-introduction inference to all these derivations followed by an $\bigvee$-introduction inference:

$$\bigvee\ \frac{\bigwedge\ \dfrac{\neg A_1, \ldots, \neg A_{i-1}, A_i, \varphi_j \quad , j < J}{\neg A_1, \ldots, \neg A_{i-1}, A_i, \bigwedge_{j<J}\varphi_j}}{\neg A_1, \ldots, \neg A_{i-1}, A_i}$$

The tree-size of this derivation is $O(S^3)$.

The final step is to combine the cedents $\neg A_1, \ldots, \neg A_{i-1}, A_i$ with $L-1$ many cuts to produce the desired refutation. Observe that for $i = L$, the Claim implies that we can derive the cedent $\neg A_1, \ldots, \neg A_{L-1}$ as $\Gamma_L = \emptyset$.

$$\text{cut}\ \frac{\neg A_1, \ldots, \neg A_{L-1} \qquad \neg A_1, \ldots, \neg A_{L-2}, A_{L-1}}{\neg A_1, \ldots, \neg A_{L-2}}$$

$$\vdots$$

$$\text{cut}\ \frac{\text{cut}\ \dfrac{\neg A_1, \neg A_2 \qquad \neg A_1, A_2}{\neg A_1} \qquad A_1}{\emptyset}$$

As each A_j is a Θ^S_{d+1}-formula, this derivation is a Θ^S_{d+1}-PK refutation of $\mathcal{A}$ of tree-size $O(S^4)$, as required. $\square$

Proof of Theorem 4. To prove (1) and (2) are equivalent, use

Proposition 11 shows	(1) for d-PK	$\Longrightarrow$	(2) for (d+1)-PK	
[5, Lemma 6] shows	(2) for (d+1)-PK	$\Longrightarrow$	(1) for d-PK	□

Although we are not able to prove the "Former Statement 5" from our previous paper [5], the results obtained in the present paper allow us to relate d-PK sequence-size proofs and (d+1)-PK tree-size proofs to (d+2)-PK proofs of constant critical cut-rank. This is useful for applications to Bounded Arithmetic: An occurrence of an induction inference in a first-order derivation in a Bounded Arithmetic theory can be unwound into a series of cuts; this can be done either with a balanced tree of height logarithmic in the length of the induction, or in a linear fashion which increases the critical cut-rank of the propositional derivation by 1. Using the latter method means that unwinding all the induction inferences in a first-order Bounded Arithmetic proof yields propositional proofs with constant critical cut-rank. Theorem 8 then implies the existence of polynomial size tree- or sequence-sized PK derivations of low formula depth. This last part is stated precisely in Corollary 12:

Corollary 12. *Let $d \in \frac{1}{2}\mathbb{N}$. For $n \in \mathbb{N}$, let $\mathcal{A}_n$ be a collection of Θ_d-cedents. Then the following conditions* (1), (2) *and* (5) *are equivalent:*

(1) $\mathcal{A}_n$ *has a* d-PK *refutation of sequence-size polynomial in n, for all n.*

(2) $\mathcal{A}_n$ *has a* (d+1)-PK *refutation of tree-size polynomial in n, for all n.*

(5) $\mathcal{A}_n$ *has a* (d+2)-PK *refutation which simultaneously has tree-size polynomial in n and constant critical cut-rank, for all n.*

The implication from (5) *to* (2) *holds more generally: namely,* (5′) *implies* (3)*:*

(5′) $\mathcal{A}_n$ *has a* (d+1)-PK *refutation which simultaneously has tree-size polynomial in n and constant critical cut-rank, for all n.*

(3) $\mathcal{A}_n$ *has a* d-PK *refutation of tree-size polynomial in n, for all n.*

Proof. Theorem 8 immediately shows that (5) implies (2), and that (5′) implies (3). The implication (2) implies (1) is part of Theorem 4.

The direction (1) implies (5) follows from a modification of the proof of Lemma 4 in [5]. We will give a sketch of the modification of that proof. A refutation of $\mathcal{A}$ of sequence-size S can be written as a sequence of cedents $\Gamma_1, \ldots, \Gamma_L{=}\emptyset$ where each Γ_i is a logical axiom, is a nonlogical axiom Γ or $\bigvee^{\mathrm{coll}}\Gamma$ with $\Gamma \in \mathcal{A}$, or is formed from previous cedents $\Gamma_1, \ldots, \Gamma_{i-1}$ by applying one of the rules of PK. Let γ_i be $\bigwedge_{j<i}\bigvee^{\mathrm{coll}}\Gamma_j$. The proof of Lemma 4 in [5] shows that the cedent $\neg\gamma_i, \gamma_{i+1}$ has a cut-free PK derivation from $\mathcal{A}$ of size $O(S^3)$. Combining these with cuts in a linear fashion from right-to-left, starting from cedent $\neg\gamma_{L-1}, \gamma_L$, yields a ($d$+2)-PK refutation of $\mathcal{A}$ of simultaneous tree-size $O(S^4)$ and critical cut-rank 1. □

It is open whether (3) implies (5′) when d equals 0 or $\frac{1}{2}$. For larger d, the implication is already given from (2) implies (5).

References

[1] M. Ajtai. The complexity of the pigeonhole principle. In *Proceedings of the 29-th Annual IEEE Symposium on Foundations of Computer Science*, pages 346–355, 1988.

[2] David A. Mix Barrington. Quasipolynomial size circuit classes. In *Proc. 7th Structure in Complexity Conference*, pages 86–93, 1992.

[3] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. *Proceedings of the London Mathematical Society*, 73(3):1–26, 1996.

[4] Paul Beame and Toniann Pitassi. An exponential separation between the parity principle and the pigeonhole principle. *Annals of Pure and Applied Logic*, 80(3):195–228, 1996.

[5] Arnold Beckmann and Samuel R. Buss. Separation results for the size of constant-depth propositional proofs. *Annals of Pure and Applied Logic*, 136:30–55, 2005.

[6] Jan Krajíček. Lower bounds to the size of constant-depth propositional proofs. *Journal of Symbolic Logic*, 59:73–86, 1994.

[7] Jan Krajíček. Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *Journal of Symbolic Logic*, 62:457–486, 1997.

[8] Jan Krajíček and Pavel Pudlák. Quantified propositional calculi and fragments of bounded arithmetic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 36:29–46, 1990.

[9] Jan Krajíček, Pavel Pudlák, and Alan Woods. Exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures and Algorithms*, 7:15–39, 1995.

[10] Jeff B. Paris and Alex J. Wilkie. Δ_0 sets and induction. In W. Guzicki, W. Marek, A. Pelc, and C. Rauszer, editors, *Open Days in Model Theory and Set Theory*, pages 237–248, 1981.

[11] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3:97–140, 1993.

[12] Alexander A. Razborov. On provably disjoint NP-pairs. Technical Report TR94-006, Electronic Colloquium in Computational Complexity (ECCC), 1994. Also available as BRIC technical report RS-94-36.