EXAMINING FRAGMENTS OF THE QUANTIFIED PROPOSITIONAL CALCULUS

STEVEN PERRON†

Abstract. When restricted to proving Σ_i^q formulas, the quantified propositional proof system G_i^* is closely related to the Σ_i^b theorems of Buss's theory S_2^i. Namely, G_i^* has polynomial-size proofs of the translations of theorems of S_2^i, and S_2^i proves that G_i^* is sound. However, little is known about G_i^* when proving more complex formulas. In this paper, we prove a witnessing theorem for G_i^* similar in style to the KPT witnessing theorem for T_2^i. This witnessing theorem is then used to show that S_2^i proves G_i^* is sound with respect to Σ_{i+1}^q formulas. Note that unless the polynomial-time hierarchy collapses S_2^i is the weakest theory in the S_2 hierarchy for which this is true. The witnessing theorem is also used to show that G_1^* is p-equivalent to a quantified version of extended-Frege for prenex formulas. This is followed by a proof that G_i p-simulates G_{i+1}^*. We finish by proving that S_2 can be axiomatized by S_2^1 plus axioms stating that the cut-free version of G_0^* is sound. All together this shows that the connection between G_i^* and S_2^i does not extend to more complex formulas.

§1. Introduction. In [9], Krajicek and Pudlak introduced the quantified propositional proof system G and its fragments. These fragments have close connections with bounded arithmetic and computational complexity. In particular, the collapse of the polynomial-time hierarchy, the bounded arithmetic hierarchy S_2, and the fragments of G are all related [9, 8, 7, 11]. Even with these close connections to important open problems in logic and computer science, little work has been done investigating the fragments of G. In this paper, we take a closer look at them.

The proof system G_i^* has informally been described as the non-uniform version of S_2^i. This is often expressed by describing the close connection between the Σ_i^b theorems of S_2^i and G_i^* proofs of Σ_i^q formulas [7]. The same type of connection exists between the theory PV and extended-Frege [2], and the theory T_2^i and G_i. In this paper, we compare these proof systems to each other and the theories to verify the accuracy of these informal descriptions.

Following Morioka, the proof system G_i^* is defined by restricting G to treelike proofs where all cut formulas are Σ_i^q [4, 12]. Note that originally G_i^* was defined by restricting all formulas, not just cut formulas, to Σ_i^q formulas [9, 7]. Informally, we can think of G_i^* as reasoning with lemmas that can be described as predicates in the ith level of the polynomial-time hierarchy.

We examine G_1^* by comparing it to extended-Frege directly. In [7], it was shown that treelike extended-Frege is p-equivalent to G_1^* with respect to quantifier-free

†This work was funded by the Natural Sciences and Engineering Research Council of Canada.

formulas. This means that, when proving quantifier-free formulas, G_1^* only needs to cut quantifier-free formulas and extension cedents. This raises the question of whether or not this holds when G_1^* is used to prove more complicated formulas. We define a quantified version of extended-Frege called GPV^*, and prove that GPV^* and G_1^* are p-equivalent with respect to all prenex formulas. This result is surprising because the class of formulas that GPV^* can cut is much less expressive than the class of formulas that G_1^* can cut. As well, this result does not fit with the view that GPV^* corresponds to PV and G_1^* with S_2^1.

We also take a look at G_i and G_{i+1}^*. If we used the connections with bounded arithmetic as a guide, we would expect G_{i+1}^* to be a strictly stronger proof system than G_i. However, in [13], Nguyen showed that this is probably not the case. This was done by showing that, under an appropriate complexity assumption, G_{i+1}^* does not simulate G_i or even cut-free G for Σ_{i+2}^q formulas. This is in contrast to a result that shows that G_{i+1}^* p-simulates G_i for Σ_{i+1}^q formulas. In this paper, we prove that, in fact, G_i is stronger than G_{i+1}^*, which is surprising. This is done by showing that G_i p-simulates G_{i+1}^* for all formulas, not just Σ_i^q formulas as in [7].

Another way of examining G_i^* is to find the weakest fragment of S_2 that can prove that G_i^* is sound. So, we are looking for a theory that proves that, if there is a G_i^* proof of a formula, then that formula is valid. Informally, this gives an upper bound on the reasoning power of G_i^*. This type of question first appeared in [2], where Cook showed that PV proves that extended-Frege is sound. This kind of result played an important role in establishing the connection between the collapse of S_2 and G [9].

In [7], it was shown that S_2^1 proves that G_1^* is sound with respect to Σ_1^q formulas. However, in [12], Morioka proved that, assuming the polynomial hierarchy does not collapse, S_2^1 does not prove that G_1^* is sound with respect to Σ_3^q formulas. This does not fit with the view that G_1^* is the non-uniform version of S_2^1. In fact, it seems that, as the quantifier complexity of the formulas we are proving grows, the reasoning power of G_1^* grows beyond any finite level of the S_2 hierarchy. For the same proof also shows that, assuming the polynomial-time hierarchy does not collapse, T_2^i does not prove that G_1^* is sound with respect to Σ_{i+2}^q formulas; however, we show that S_2^{i+1} does. In fact, we show that S_2^{i+1} proves G_{i+1}^* is sound with respect to Σ_{i+2}^q formulas. Informally this means that the reasoning power of G_1^* relative to Σ_{i+2}^q formulas is not stronger than the reasoning power of S_{i+1}^2.

This leads to the final way of examining G_1^*. In [9], Krajicek and Pudlak were able to prove that S_2 can be axiomatized by S_2^1 plus axioms stating G_i is sound relative to Σ_i^q formulas, for $i \in \mathbb{N}$. We show that the same is true when G_i is replaced by G_1^*. In fact, we can replace G_i by the cut-free version of G^*. This is interesting because it confirms that the reasoning power of G_1^* is not closely related to any finite level of S_2, but, in some sense, it captures the reasoning power of all of S_2.

The main tool used to prove some of these theorems is a witnessing theorem in the style of the KPT witnessing theorem [8]. The original KPT witnessing theorem describes how hard it is to witness Σ_{i+3}^b theorems of T_2^i, for $i > 0$.

It also holds for $i = 0$ with PV in place of T_2^i. This theorem has been used to prove that the collapse of the S_2 hierarchy implies the collapse of the polynomial-time hierarchy [8], and to show that certain weak theories do not prove the Σ_1^b replacement scheme, relative to some complexity assumptions [6]. In this paper, we adapt the statement of the KPT witnessing theorem to G_i^*, and then prove it. The main difficulty is that proofs of the KPT witnessing theorem rely on the cut-elimination theorem, which unfortunately causes the size of the proof to increase exponentially. We must avoid this increase, so we have to find a way to work around cut formulas.

The paper is organized as follows. In Section 2, we give the basic definitions and notations. Note that we will be using two-sorted theories of bounded arithmetic (V^i) in place of the single-sorted theories (S_2^i). In Section 3, we prove the witnessing theorem for G_i^*. In Section 4, we use the witnessing theorem to prove that GPV^* p-simulates G_1^*. In Section 5, we show that G_i p-simulates G_{i+1}^*. In Section 6, we prove the Σ_{i+1}^q reflection principle for G_i^* in S_2^i. In Section 7, we give a new axiomatization of S_2.

I should mention that this paper is an expanded version of a conference paper [14]. The main difference is this paper has the witnessing theorem for G_i^* proofs of any formula as opposed to G_1^* proofs of prenex formulas. At this point, I would like to thank my supervisor Stephen Cook for comments on earlier versions of this paper.

§2. Basic Definitions And Notation.

2.1. Two-Sorted Bounded Arithmetic. In the introduction, the results were stated for the theories S_2^i. However, we will use two-sorted theories of bounded arithmetic. We follow the presentation in [3, 5]. The two sorts are numbers and binary strings (aka finite sets). The numbers are intended to range over the natural numbers and will be denoted by lower-case letters. For example, i, j, x, y, and z will often be used for number variables; r, s, and t will be used for number terms; and f, g and h will be used for functions that return numbers. The sets are intended to be finite sets of natural numbers. Since the sets are finite, they can be coded by binary strings where the ith bit is 1 if i is in the set. The strings will be denoted by upper- case letters. The letters X,Y, and Z will often be used for string variables; and F, G and H will be used for functions that return strings.

The base language is

$$\mathcal{L}_A^2 = \{0, 1, +, \times, <, =, =_2, \in, ||\}\,.$$

The constants 0 and 1 are number constants. The functions $+$ and $\times$ take two numbers as input and return a number–the intended meanings are the obvious ones. The language also includes two binary predicates that take two numbers: $<$ and $=$. The predicate $=_2$ is meant to be equality between strings, instead of numbers. In practice, the 2 will not be written because which equality is meant is obvious from the context. The membership predicate $\in$ takes a number i and a string X. It is meant to be true if the ith bit of X is 1 (or i is in the set X). This will also be written as $X(i)$. The final function $|X|$ takes a string as input and returns a number. It is intended to be the number of bits needed to write

X when leading zeros are removed (or the least upper bound of the set X). The set of axioms 2BASIC is the set of defining axioms for $\mathcal{L}_A^2$.

We use $\exists X < b\ \phi$ as shorthand for $\exists X[(|X| < b) \wedge \phi]$. The shorthand $\forall X < b\ \phi$ means $\forall X[(|X| < b) \supset \phi]$. The set $\Sigma_0^B = \Pi_0^B$ is the set of formulas whose only quantifiers are bounded number quantifiers. For $i > 0$, the set Σ_i^B is the set of formulas of the form $\exists \vec{X} < \vec{t}\phi$ where ϕ is a Π_{i-1}^B formula. For $i > 0$, the set Π_i^B is the set of formulas of the form $\forall \vec{X} < \vec{t}\phi$ where ϕ is a Σ_{i-1}^B formula.

Now we can define the two main axiom schemes:

$$\Sigma_i^B\text{-comp: } \exists X \leq b \forall i < b[X(i) \leftrightarrow \phi(i)],$$
$$\Sigma_i^B\text{-string-ind: } [\phi(\emptyset) \wedge \forall X[\phi(X) \supset \phi(S(X))]] \supset \phi(Y)$$

where $\phi(i)$ is a Σ_i^B formula, and, for Σ_i^B-COMP, ϕ does not contain X, but may contain other free variables. The constant $\emptyset$ is the empty string, and the function $S(X)$ interprets X as a binary number and adds 1 to it. Note that we still view Σ_i^B-string-ind as a formula over $\mathcal{L}_A^2$. We simply replace the instances of $\emptyset$ and $S(X)$ by their Σ_0^B bit-definition.

We can now define two hierarchies of theories.

Definition 2.1. *The theory V^i is axiomatized by the 2BASIC axioms plus Σ_i^B-comp. The theory TV^i is axiomatized by the 2BASIC axioms, Σ_0^B-comp, and Σ_i^B-string-ind.*

For $i > 0$, V^i corresponds to S_2^i, and TV^i corresponds to T_2^i in that they are RSUV-isomorphic [5].

Another theory we often use is VPV, a universal theory with a function symbol for every polynomial-time function. The function symbols have the following defining axioms based on Cobham's Theorem:

Definition 2.2 (PV function symbols). *The language $\mathcal{L}_{FP}$ is the smallest set satisfying the following:*

1. *$\mathcal{L}_{FP}$ includes $\mathcal{L}_A^2 \cup \{pd, CHOP\}$.*
2. *For each open formula $\phi(z, \vec{x}, \vec{X})$ over $\mathcal{L}_{FP}$ and term $t = t(\vec{x}, \vec{X})$ over $\mathcal{L}_A^2$, there is a string function $F_{\phi,t}$ and a number function $f_{\phi,t}$ in $\mathcal{L}_{FP}$.*
3. *For each triple G, H, t, where $G(\vec{x}, \vec{X})$ and $H(y, \vec{x}, \vec{X}, Z)$ are functions in $\mathcal{L}_{FP}$ and $t(y, \vec{x}, \vec{X})$ is an $\mathcal{L}_A^2$ term, there is a function $F_{G,H,t}$ in $\mathcal{L}_{FP}$.*

The 2BASIC axiom define the function symbols in $\mathcal{L}_A^2$. The rest of the functions symbols have the following defining axioms:

- $pd(0) = 0$, $x \neq 0 \supset pd(x) + 1 = x$
- $CHOP(X, y)(i) \leftrightarrow i < y \wedge X(i)$
- $F_{\phi,t}(\vec{x}, \vec{X})(i) \leftrightarrow i < t(\vec{x}, \vec{X}) \wedge \phi(i, \vec{x}, \vec{X})$
- $i < f_{\phi,t}(\vec{x}, \vec{X}) \supset \neg\phi(i, \vec{x}, \vec{X})$
- $f_{\phi,t}(\vec{x}, \vec{X}) < t(\vec{x}, \vec{X}) \supset \phi(f_{\phi,t}(\vec{x}, \vec{X}), \vec{x}, \vec{X})$
- $F_{G,H,t}(0, \vec{x}, \vec{X}) = G(0, \vec{x}, \vec{X})$
- $F_{G,H,t}(y + 1, \vec{x}, \vec{X}) = CHOP(H(y, \vec{x}, \vec{X}, F_{G,H,t}(y, \vec{x}, \vec{X})), t(y, \vec{x}, \vec{X}))$

The theory VPV is axiomatized by quantifier-free equivalents of the 2BASIC axioms, induction on all open $\Sigma_0^B(PV)$ formulas, and the defining axioms for all of the $\mathcal{L}_{FP}$ function symbols. See [3, 5] for more information on VPV.

Another scheme of formulas we use is the Σ_i^B-MAX scheme:

$$\exists x < b\phi(x) \supset \exists x < b[\phi(x) \wedge \forall y < b(x < y \supset \neg\phi(y))]$$

where ϕ is Σ_i^B. This scheme essentially says that, if there exists a value for x less than b that satisfies $\phi(x)$, then there exists a maximum x less than b that satisfies $\phi(x)$. It can be shown that Σ_i^B-MAX is provable in V^i ([5], Corollary 5.8).

From time to time, we will use functions symbols that are not in $\mathcal{L}_A^2$. The first is $X(i,j) \equiv X(\langle i,j\rangle)$, where $\langle i,j\rangle = (i+j)(i+j+1)+2j$ is the pairing function. It can be thought of as a two dimensional array of bits. The second is the row function. The notation we use is $X^{[i]}$. This functions returns the ith row of the two dimensional array X. In the same way, we can also describe three dimensional arrays. We also want to pair string. So if $X = \langle Y_1, Y_2\rangle$, then $X^{[0]} = Y_1$ and $X^{[1]} = Y_2$. Note that, if we add these functions with their Σ_0^B defining axioms to the theory V^i, we get a conservative extension. They can also be used in the induction axioms [3]. This means that, if there is a V^i proof of a formula that uses these functions, there is a V^i proof of the same formula that does not use these functions.

2.2. Quantified Propositional Calculus. We are also interested in quantified propositional proof systems. The proof systems we use were originally defined in [9], and then they were redefined in [4, 12], which is the presentation we follow.

The set of connectives are $\{\wedge, \vee, \neg, \exists, \forall, \top, \bot\}$, where $\top$ and $\bot$ are constants for true and false, respectively. Formulas are built using these connectives in the usual way. We will often refer to formulas by the number of quantifier alternations.

DEFINITION 2.3. *The set of formulas $\Sigma_0^q = \Pi_0^q$ is the set of quantifier-free propositional formulas. For $i > 0$, the set of Σ_i^q (Π_i^q) formulas is the smallest set of formulas that contains Π_{i-1}^q (Σ_{i-1}^q) and is closed under $\wedge$, $\vee$, existential (universal) quantification, and if $A \in \Pi_i^q$ ($A \in \Sigma_i^q$) then $\neg A \in \Sigma_i^q$ ($\neg A \in \Pi_i^q$).*

The first proof system, from which all others will be defined, is the proof system G. This proof system is a sequent calculus based on Gentzen's system LK. The system G is essentially the DAG-like, propositional version of LK. We will not give all of the rules, but will mention a few of special interest.

The cut rule is

$$\text{cut } \frac{A, \Gamma \to \Delta \qquad \Gamma \to \Delta, A}{\Gamma \to \Delta}$$

In this rule, we call A the cut formula. There are also four rules that introduce quantifiers:

$$\exists\text{-left } \frac{A(x), \Gamma \to \Delta}{\exists z A(z), \Gamma \to \Delta} \qquad \exists\text{-right } \frac{\Gamma \to \Delta, A(B)}{\Gamma \to \Delta, \exists z A(z)}$$

$$\forall\text{-left}\ \frac{\Gamma \to \Delta, A(x)}{\Gamma \to \Delta, \forall z A(z)} \qquad \forall\text{-right}\ \frac{A(B), \Gamma \to \Delta}{\forall z A(z), \Gamma \to \Delta}$$

These rules have conditions on them. In $\exists$-left and $\forall$-right, the variable x must not appear in the bottom sequent. In these rules, x is called the eigenvariable. In the other two rules, the formula B must be a Σ_0^q formula, and no variable that appears free in B can be bound in $A(x)$.

The initial sequents of G are sequents of the form $\to \top$, $\bot \to$, or $x \to x$, where x is any propositional variable. A G proof is a series of sequents such that each sequent is either an initial sequent or can be derived from previous sequents using one of the rules of inference. The proof system G_i is G with cut formulas restricted to Σ_i^q formulas.

We define G^* as the treelike version of G. So, a G^* proof is a G proof where each sequent in used as an upper sequent in an inference at most once. A G_i^* proof is a G^* proof in which cut formulas are prenex Σ_i^q. In [12], it was shown that, for treelike proofs, it did not matter if the cut formulas in G_i^* were prenex or not. So when we construct G_i^* proofs the cut formulas will not always be prenex, but that does not matter.

To make proofs simpler, we assume that all treelike proofs are in *free-variable normal form.*

Definition 2.4. *A parameter variable for a G_i^* proof π is a variable that appears free in the final sequent of π. A proof π is in* free-variable normal form *if (1) every non-parameter variable is used as an eigenvariable exactly once in π, and (2) parameter variables are not used as eigenvariables.*

Note that, if a proof is treelike, we can always put it in free-variable normal form by simply renaming variables. In fact, VPV proves that every treelike proof can be put in free-variable normal form.

A useful property of these proof systems is the *subformula property.* It can be shown in VPV that every formula in a G_i^* proof is an ancestor (and therefore a subformula) of a cut formula or a formula in the final sequent. This is useful because it tells us that any non-Σ_i^q formula in a G_i^* proof must be an ancestor of a final formula.

2.3. Truth Definitions. In order to reason about the proof systems in the theories, we must be able to reason about quantified propositional formulas. We follow the presentation in [7, 9].

Formally formulas will be coded as string, but we will not distinguish between a formula and its encoding. So if F is a formula, we will use F as the string encoding the formula as well. The method of coding a formula can be found in [4]. The encoding of an assignment A will be a set of pairs $\langle i, 0\rangle$ and $\langle i, 1\rangle$ which mean that the variable x_i is assigned false and true, respectively.

The truth definition will be defined in the usual way. For Σ_i^q formulas, we will construct a formula that essentially says there exists an evaluation of the formula and it evaluates to true. For Π_i^q formula, the statement is the same except we say that all evaluations evaluate to true. The definition will be given recursively. Because of the potentially large and repetitive nature of the definition, we will only give part of the definition and leave it to the reader to complete it.

Given a Σ_i^q formula F, an evaluation of that formula will be a series of lines. Each line will consist of a truth value and a subformula of F. Plus each line will have to be consistent with previous lines. With this in mind, we have the following definition.

DEFINITION 2.5. *We recursively define $A \models_i F$. If F is a Σ_i^q formula, then*

$$A \models_i F \equiv \exists X \exists E \; eval_i^{\exists}(E, A, X, F) \wedge \exists n < |E|, E^{[n]} = \langle \top, F \rangle,$$

where $eval_i^{\exists}(E, A, X, F)$ is a formula saying that E is a series of lines assigning truth values to subformulas of F and X assigns values to the outermost even existential quantifiers and the outermost odd universal quantifiers. An even quantifier is one that is in the scope of an even number of $\neg$, and an odd quantifier is on that is in the scope of an odd number of $\neg$. More formally $eval_i^{\exists}(E, A, X, F)$ is the conjunction of the following. Note that we do not give the bounds on the quantified variables, but the reader can fill in what they should be.

- *For every line l, if the outermost connective of the formula is $\wedge$ and the formula is true, then there are earlier lines j_1, j_2 saying both left subformula F_1 and right subformula F_2 are true.*

  $$\forall l \exists F_1 \exists F_2 \exists j_1 \exists j_2,$$
  $$E^{[l]} = \langle \top, F_1 \wedge F_2 \rangle \supset E^{[j_1]} = \langle \top, F_1 \rangle \wedge E^{[j_2]} = \langle \top, F_2 \rangle$$

  Note that the case of $\vee$ with a false formula is handled the same way.
- *For every line l, if the outermost connective of the formula is $\wedge$ and the formula is false, then there is an earlier line j saying one of the left subformula F_1 or right subformula F_2 is false.*

  $$\forall l \exists F_1 \exists F_2 \exists j,$$
  $$E^{[l]} = \langle \bot, F_1 \vee F_2 \rangle \supset E^{[j]} = \langle \bot, F_1 \rangle \vee E^{[j]} = \langle \bot, F_2 \rangle$$

  Note that the case of $\vee$ with a true formula is handled the same way.
- *For every line l, if the outermost connective of the formula is $\neg$ and the formula is true, then there is a previous line j saying the subformula is false.*

  $$\forall l \exists F_1 \exists j,$$
  $$E^{[l]} = \langle \top, \neg F_1 \rangle \supset E^{[j]} = \langle \bot, F_1 \rangle$$

  Note that this is the only case where the truth value changes, so the truth value can also be viewed as the parity of the number of negations that were passed to reach this subformula.
- *For every line l, if the outermost connective of the formula is $\exists$ and the formula is true, then there is a previous line j with a witness for the quantifier and X gives us that value.*

  $$\forall l \exists F_1 \exists j,$$
  $$E^{[l]} = \langle \top, \exists x_n F_1(x_n) \rangle \supset (E^{[j]} = \langle \top, F_1(x_n) \rangle \wedge (\langle n, 0 \rangle \in X \vee \langle n, 1 \rangle \in X))$$

  Note that the case of $\forall$ with a false formula is handled the same way.

- *For every line l, if the outermost connective of the formula is $\exists$ and the formula is false, then the formula is a Σ^q_{i-1} formula, and it is false according to $\models_{i-1}$.*

  $\forall l \exists F_1,$

  $E^{[l]} = \langle \bot, \exists y_n F_1(y_n) \rangle \supset \exists y_n F_1(y_n) \in \Sigma^q_{i-1} \wedge (X \cup A) \models_{i-1} \exists y_n F_1(y_n)$

  Note that the case of $\forall$ with a true formula is handled the same way.
- *For every line l, if the formula is a single variable, then the truth value is consistent with A.*

$$\begin{aligned} &\forall l, \\ &E^{[l]} = \langle \top, x_n \rangle \supset \langle x_n, 1 \rangle \in A \\ \wedge &E^{[l]} = \langle \bot, x_n \rangle \supset \langle x_n, 0 \rangle \in A \end{aligned}$$

If F is a Π^q_i formula, then

$$A \models_i F \equiv \forall Y \forall E \; eval^{\forall}_i(E, A, Y, F) \supset \exists n E^{[n]} = \langle \top, F \rangle$$

where $eval^{\forall}_i(E, A, Y, F)$ is almost the same as $eval^{\exists}_i$ except Y now gives a truth value for the even universally-quantified variables and the odd existentially-quantified variables.

Notice that in $eval^{\exists}_i$ if the outermost connective is $\forall$ and we want to falsify it, then it is treated like $\exists$. The connectives $\wedge$ and $\vee$ are also treated the same when we are trying to satisfy one and falsify the other. When we see a $\forall$ and we want to satisfy the formula, we know the quantifier complexity of the formula has dropped. Therefore, we can get the value of this formula recursively. If we are looking at a Σ^q_0 formula the recursive case never comes up.

For a Σ^q_i formula, we are saying there is an evaluation of the formula that says it is true. For a Π^q_i formula, we are saying that all evaluations of the formula say it is true. This is an important difference since a Σ^q_i formula is false if there is no evaluation of the formula, but a Π^q_i formula would be true.

For $i > 0$, this gives a Σ^B_i definition for $A \models_i F$ and, for $i = 0$, it has a $\Sigma^B_0(PV)$ definition in VPV. If F is a Π^q_i formula, the definition is Π^B_i.

Given a formula $F \equiv \bigwedge_{i=0}^{n} F_i$, there is a PV function $Parse_{\wedge}(F, j)$ that outputs $F_{min(j,n)}$. The same goes for $\vee$ in place of $\wedge$. The theory VPV proves the Tarski conditions for the truth definition.

LEMMA 2.6 (Tarski's Conditions). *VPV proves the following*

1. $(A \models_i F_1 \wedge F_2) \leftrightarrow (A \models_i F_1 \wedge A \models_i F_2)$
2. $(A \models_i F_1 \vee F_2) \leftrightarrow (A \models_i F_1 \vee A \models_i F_2)$
3. $(A \models_i F) \leftrightarrow (\forall j \leq |F| \; A \models_i Parse_{\wedge}(F, j))$ *(where* $F \equiv \bigwedge_{j=0}^{n} F_j$ *and* $F \in \Pi^q_i$*)*
4. $(A \models_i F) \leftrightarrow (\exists j \leq |F| \; A \models_i Parse_{\vee}(F, j))$ *(where* $F \equiv \bigvee_{j=0}^{n} F_j$ *and* $F \in \Sigma^q_i$*)*
5. $(A \models_i \neg F) \leftrightarrow \neg(A \models_i F)$
6. $(A \models_i \exists \vec{x} F(\vec{x})) \leftrightarrow \exists X (A \cup X \models_i F(\vec{x}))$ *(for* $F \in \Sigma^q_i$*)*
7. $(A \models_i \forall \vec{x} F(\vec{x})) \leftrightarrow \forall X (A \cup X \models_i F(\vec{x}))$ *(for* $F \in \Pi^q_i$*)*

8. $(A \models_i F) \leftrightarrow (A \models_{i-1} F)$ *(for* $F \in \Sigma^q_{i-1} \cup \Pi^q_{i-1}$*).*

PROOF. (1) Suppose $A \models_i F_1 \wedge F_2$ and the formula is Σ^q_i, then there is an evaluation of this formula. This evaluation would contain the line $(\top, F_1 \wedge F_2)$. Therefore this evaluation would also contain lines of the form $(\top, F_1)$ and $(\top, F_2)$. This means we have evaluations of F_1 and F_2. Suppose $A \models F_1 \wedge A \models F_2$. Then there exist evaluations of these formulas. An evaluation for $F_1 \wedge F_2$ is obtained from these evaluations by combining these evaluations and adding a new line using Σ^B_0-COMP. If the formula is Π^q_i, the proof is similar.

(2) The same way as (1).

(3) Suppose $A \models_i F$ is false and $\forall j \leq |F|$ $A \models_i Parse_\wedge(F, j)$ is true, where $F \equiv \bigwedge_{j=0}^{n} F_j$. Let F^m be the formula $\bigwedge_{j=0}^{m} F_j$. This means $F^m \equiv F^{m-1} \wedge F_m$. By the first assumption, there is an evaluation of $F \equiv F^n$ with a line $\langle \bot, F \rangle$. If there is a line in the evaluation of the form $\langle \bot, F^{m+1} \rangle$, then there is a line $\langle \bot, F^m \rangle$. Note that $\langle \bot, F_m \rangle$ cannot appear in a line by the second assumption. So it follows by induction that there is a line saying that $F^1 \equiv F_1$ is false, but this contradicts the second assumption.

Now suppose $\forall j \leq |F|$ $A \models_i Parse_\wedge(F, j)$ is false. Then there exists an evaluation of F_j, for some j, with a line $\langle \bot, F_j \rangle$. To this evaluation we can append the lines $\langle \bot, F^m \rangle$ for $j \leq m \leq n$ using Σ^B_0-COMP. This shows that $A \models_i F$ is false.

Note that the proof of (3) does not work if F is a Σ^q_i formula since $A \models_i F$ could be false because there is no evaluation of F. There is not necessarily an evaluation that shows F is false.

(4) This is the dual of (3).

(5) We assume F is a Π^q_i formula. The case where F is a Σ^q_i formula is essentially the same. Suppose $(A \models_i \neg F)$. Then there exists an evaluation of F with a line $\langle \top, \neg F \rangle$. This means the evaluation also has a line $\langle \bot, F \rangle$, proving $\neg(A \models_i F)$. Suppose $\neg(A \models_i F)$. There there exists an evaluation of F with line $\langle \bot, F \rangle$. The line $\langle \top, \neg F \rangle$ can be appended to this evaluation, proving $(A \models_i \neg F)$.

(6) This follows directly from the $\exists X$ in the definition of $\models_i$.

(7) This is the dual of (6).

(8) This follows directly from the recursive nature of the definition. ⊣

Valid formulas (or tautologies) are defined as

$$TAUT_i(F) \equiv \forall A, (\text{“A is an assignment to the variables of } F\text{”} \supset A \models_i F)$$

This truth definition can be extended to define the truth of a sequent. So, if $\Gamma \to \Delta$ is a sequent of $\Sigma^q_i \cup \Pi^q_i$ formulas, then

$$\begin{aligned}(A \models_i \ \Gamma \to \Delta) \equiv &\text{“there exists a formula in } \Gamma \text{ that } A \text{ does not satisfy”} \\ &\vee \ \text{“there exists a formula in } \Delta \text{ that } A \text{ satisfies”}\end{aligned}$$

Another important formula we will use is the reflection principle for a proof system. We define the Σ^q_i reflection principle for a proof system P as

$$\Sigma^q_i\text{-RFN}(P) \equiv \forall F \forall \pi, (\text{“}\pi \text{ is a } P \text{ proof of } F\text{”} \wedge F \in \Sigma^q_i) \supset TAUT_i(F)$$

This formula essentially says that, if there exists a P proof of a Σ_i^q formula F, then F is valid. Another way of putting it is to say that P is sound when proving Σ_i^q formulas.

§3. KPT Style Witnessing for Fragments of G. In bounded arithmetic, a useful tool has been the KPT witnessing theorem [8]. In the simplest case, the KPT witnessing theorem describes how to witness the Σ_2^B theorems of VPV. The original theorem was more general, but we state it here for the simplest case.

THEOREM 3.1 (KPT Witnessing [8]). *Suppose $VPV \vdash \forall X \exists Y \forall Z \phi(X, Y, Z)$, where ϕ is a Σ_0^B formula. Then there exists a finite sequence of PV function symbols $F_1, F_2, \ldots, F_k$ such that*

$$\begin{aligned} VPV \vdash \forall X \forall W \ & \phi(X, F_1(X), W^{[1]}) \\ & \vee \phi(X, F_2(X, W^{[1]}), W^{[2]}) \\ & \vdots \\ & \vee \phi(X, F_k(X, W^{[1]}, W^{[2]}, \ldots, W^{[k-1]}), W^{[k]}) \end{aligned}$$

Informally, this can be viewed as an interactive computation between a student, who runs in polynomial time, and an all-knowing teacher. Given a value for X, the student's goal is the find a witness for $\exists Y \forall Z \phi(X, Y, Z)$. The student starts by computing $F_1(X)$. If that is not a witness, the teacher responds with a counter example $W^{[1]}$. Using that the students makes a second guess by computing F_2. The teacher responds with $W^{[2]}$, and this process continues.

Our goal is to get a similar theorem for G_1^*, and to extend this to G_i^*. The rest of this section is organized as follows. We start by stating the analog of the above theorem for G_1^*. Using this as a starting point, we then define the concepts needed to prove this theorem. Our presentation will be based on a proof of the Herbrand Theorem. We then prove an analog of the Herbrand Theorem for G_1^*, and as a corollary we get a proof of the KPT Witnessing Theorem for G_1^*. In the second subsection, we explain how to generalize the Herbrand Theorem for G_1^* so it works for G_i^*.

3.1. Witnessing for G_1^*. In adapting the KPT Witnessing Theorem for G_1^*, the first obstacle comes in the statement of the theorem. The theory VPV has access to function symbols that correspond to the polynomial-time functions, but, in G_1^*, there are no function symbols. To fix this, we use the idea of an extension cedent from [5].

DEFINITION 3.2. *An* extension cedent *is a series of formulas of the form*

$$e_1 \leftrightarrow E_1, e_2 \leftrightarrow E_2, \ldots, e_n \leftrightarrow E_n$$

such that E_i is a Σ_0^q formula that does not mention the variables $e_i, e_{i+1}, \ldots, e_n$. We say that e_i depends on a variable q if E_i mentions q or E_i mentions a variable that depends on q.

Observe that an extension cedent is really a description of a circuit, and that polynomial-size circuits are the nonuniform version of polynomial-time functions. So extension cedents replace the functions.

THEOREM 3.3 (KPT Witnessing for G_1^*). *There exists a PV (polynomial-time) function F such that VPV proves the following. Let π be a G_1^* proof of a prenex Σ_2^q formula $A(\vec{p}) \equiv \exists\vec{x}\forall\vec{y}B(\vec{x},\vec{y},\vec{p})$, where $B(\vec{x},\vec{y},\vec{p})$ is a Σ_0^q formula with all free variables shown. Then, given π, F outputs a G_0^* proof of a sequent $\Lambda \to \Theta$ where*

1. *Θ is a series of formulas of the form $B(\vec{e}^i,\vec{q}^i,\vec{p})$, where $\vec{e}^i \in E$*
2. *Λ is an extension cedent defining a new set of variables E in terms of $\vec{q}^1,\ldots,\vec{q}^n$ and $\vec{p}$,*
3. *$\vec{e}^i$ does not depend on $\vec{q}^j$, for $j \geq i$, and*
4. *$\vec{q}^i$ and $\vec{q}^j$ are disjoint.*

Before we prove this theorem, notice that this is similar to the KPT Witnessing theorem for VPV. The row $W^{[i]}$ corresponds to $\vec{q}^i$, and F_i corresponds to the circuits defining $\vec{e}^i$. The major difference is that the number of rounds in the student-teacher game is not constant; it can grow polynomially in the size of the proof.

One way of proving the KPT Witnessing Theorem is to observe that it is a corollary to the Herbrand Theorem. So the idea behind our proof is to adjust the proof-theoretic proof of the Herbrand Theorem. See [1] Section 3 for an outline of the proof we use as a model. The main difference between our proof and that proof is that cut elimination cannot be used since it causes an exponential increase in the size of the proof. To get around this problem, we use the idea in [5] to prove that extended-Frege p-simulates G_1^*. The Σ_1^q cut formulas are turned into Σ_0^q cut formulas by witnessing the existential quantifiers with extension variables.

We prove the Herbrand Theorem for all Σ_i^q formulas, but before we can state the general theorem, we need a few definitions. The first one has more to do with notation. The q variables come from the eigenvariables in the G_1^* proof. To make it easier to refer to these variables, we use the following notation:

Notation 3.4. Let π be a G^* proof. Then the set Q_π will be the set of variables that are used as eigenvariables in π. If S is a sequent in π, then $Q_{\pi,S}$ will be the set of variables that are used as eigenvariables in the subproof of π ending with S. We will refer to $Q_{\pi,S}$ as Q_S when π is understood.

Note that π is treelike, and, if it is in free-variable normal form and S is derived from S_1 and S_2, then $Q_S = Q_{S_1} \cup Q_{S_2}$, and $Q_{S_1} \cap Q_{S_2} = \emptyset$.

The general witnessing theorem will be for G_1^* proofs of any formula A. In the end, we want a G_0^* proof of a sequent $\Lambda \to A^*$, where A^* is an instance of of an $\vee$-expansion of A defined below.

From now on we assume quantifiers do not appear in the scope of a $\neg$. If we did not assume this, we would have to add a separate cases for when quantifiers appear in the scope of an odd number of quantifiers and an even number.

DEFINITION 3.5 ($\vee$-expansion). *An $\vee$-expansion of a formula A is any formula that can be obtained from A by a finite number of applications of the following rule:*

(α) *If A^* is an $\vee$-expansion of A and B is a non-Σ_0^q subformula of A^*, then replacing B by $B \vee B'$, where B' is B with renamed quantified variables, in A^* yields another $\vee$-expansion of A.*

Note that A is an $\vee$-expansion of A.

DEFINITION 3.6 ((Q,E)-instance). *Let Q and E be disjoint sets of variables. A* (Q,E)-instance *of a formula A is a quantifier-free formula A' obtained from A by replacing universally-quantified variables by distinct variables in Q and existentially-quantified variables by distinct variables in E.*

EXAMPLE 3.7. *If*

$$A \equiv B_1 \wedge \exists x[B_2(x) \wedge \forall y B_3(x,y)]$$

and $B_2(x)$ is not a Σ_0^q formula, then

$$A^* \equiv B_1 \wedge \exists x[(B_2(x) \vee B_2(x)) \wedge (\forall y B_3(x,y) \vee \forall y' B_3(x,y'))]$$

is an $\vee$-expansion of A. This can be seen by replacing $B_2(x)$ and $\forall y B_3(x,y)$. We renamed the copy of y to y' to emphasis it is now a different quantified variable. If $q_1, q_2 \in Q$ and $e \in E$, then the formula

$$B_1 \wedge [(B_2(e) \vee B_2(e)) \wedge (B_3(e,q_1) \vee B_3(e,q_2))]$$

is a (Q,E)-instance of A^*, but

$$B_1 \wedge [(B_2(e) \vee B_2(e)) \wedge (B_3(e,q_1) \vee B_3(e,q_1))]$$

is not because y and y' were replaced by the same variable.

For another example, consider a prenex formula

$$\exists \vec{x}^1 \forall \vec{y}^1 \ldots \exists \vec{x}^n \forall \vec{y}^n B(\vec{x}^1, \vec{y}^1, \ldots, \vec{x}^n, \vec{y}^n),$$

where B is a Σ_0^q formula. Then an instance of an $\vee$-expansion of this formula is a formula of the form

$$B(\vec{e}^{1,1}, \vec{q}^{1,1}, \ldots, \vec{e}^{1,n}, \vec{q}^{1,n}) \vee \ldots \vee B(\vec{e}^{m,1}, \vec{q}^{m,1}, \ldots, \vec{e}^{m,n}, \vec{q}^{m,n}).$$

So in Theorem 3.3, the disjunction of the formulas in Θ is a (Q_π, E)-instance of A. Because of this, Theorem 3.3 is simply a special case of the Herbrand Theorem for G_1^* below.

Observe that in Theorem 3.3, there is an ordering on the variables. Namely the variables $\vec{q}^i$ come before the variables $\vec{q}^{i+1}$. We could also extend this ordering to include the extension variables. An extension variable would have to be larger than every variable it depends on. For the general case, we want something similar. To make the proof simpler, we will use $\prec$ to refer to this ordering. The ordering $\prec$ orders the eigenvariables Q and the extension variables E. Then A^* will be more than a (Q,E)-instance; it will be a $(Q,E,\prec)$-instance.

DEFINITION 3.8. *$(Q, E, \prec)$-instance Let B be a (Q, E)-instance of a formula A, and let $\prec$ be an ordering on $Q \cup E$. Then B is a $(Q, E, \prec)$-instance of A if $z_1 \prec z_2$ whenever z_2 replaces a quantified variable that is in the scope of the quantified variable that z_1 replaced.*

EXAMPLE 3.9. *Take A^* from the previous example. Then*

$$B \equiv B_1 \wedge [(B_2(e) \vee B_2(e)) \wedge (B_3(e, q_1) \vee B_3(e, q_2))]$$

is a (Q, E)-instance of A^*. If B is a $(Q, E, \prec)$-instance, then we know that $e \prec q_1$ and $e \prec q_2$ since $\forall y$ and $\forall y'$ are in the scope of the $\exists x$. Note that it does not matter if $q_1 \prec q_2$ or if $q_2 \prec q_1$ since $\forall y$ is not in the scope of $\forall y'$ and vice versa.

The idea of an instance is essentially the witnessing substitution from [1]. Now we are prepared to state the general theorem.

THEOREM 3.10 (Herbrand Theorem for G_1^*). *There exists a PV function F such that VPV proves the following. Let π be a G_1^* proof of A. Then, given π, F outputs a G_0^* proof of a sequent $\Lambda \to A^*$ and a total ordering $\prec$ of the variables $Q_\pi \cup E$, where E is a set of variables that do not appear in π, with the following properties:*

- *Λ is an extension cedent defining the variables in E in terms of Q_π and the free variables of A;*
- *for $e \in E$, if e depends on a variable $p \in Q_\pi \cup E$, then $p \prec e$; and*
- *A^* is a $(Q_\pi, E, \prec)$-instance of an $\vee$-expansion of A*

PROOF. The G_0^* proof that we are looking for will be constructed by changing π one sequent at a time starting with the initial sequents and working our way down. To simplify this construction, we use the "multiplicative" form of two hypothesis rules instead of the "additive" form. For example, the multiplicative form of $\wedge$-right is

$$\frac{\Gamma_1 \to \Delta_1, A \qquad \Gamma_2 \to \Delta_2, B}{\Gamma_1, \Gamma_2 \to \Delta_1, \Delta_2, A \wedge B}$$

We use this form instead of the more standard form

$$\frac{\Gamma \to \Delta, A \qquad \Gamma \to \Delta, B}{\Gamma \to \Delta, A \wedge B}$$

This is something that was also done in [1]. The advantage of the multiplicative form is that, except for the principal formula, each formula in the bottom sequent has a single parent in the upper sequents. So, in essence, we have removed implicit contractions. We also ignore the order of the formulas in the sequents. So a sequent is a pair of multi-sets. One set for the left side of the sequent, and one set for the right side.

Let S be any sequent in π. By the subformula property of G_1^*, S is of the form

$$\Gamma \to \Delta, \Omega,$$

where Γ and Δ are possibly empty sets of Σ_1^q formulas that are not ancestors of the final formula and Ω is a possibly empty set of formulas that are ancestors of the final formula. Recall that we are assuming there are no quantifiers in the

scope of any $\neg$. We want to define a PV function that outputs a G_0^* proof of a sequent

$$S' \equiv \Lambda, \Gamma' \to \Delta', \Omega',$$

and a total ordering $\prec$ on $Q_S \cup E$ where

1. Γ' is obtained from Γ by replacing each formula $\exists \vec{z} D(\vec{z})$ by $D(\vec{q})$, where D is Σ_0^q and $\vec{q} \in Q_S$. (We use different $\vec{q}$ for different formulas.)
2. Δ' is obtained from Δ be replacing each formula $\exists \vec{z} D(\vec{z})$ by $D(\vec{e})$, where D is Σ_0^q and $\vec{e} \in E$. (We use different $\vec{e}$ for different formulas.)
3. Λ is an extension cedent defining E in terms of Q_S and the free variables of S;
4. for $e \in E$, if e depends on a variable $p \in Q_S \cup E$, then $p \prec e$;
5. Ω' is obtained from Ω by replacing each formula B by a $(Q_S, E, \prec)$-instance of an $\vee$-expansion of B; and
6. each $q \in Q_S$ appears in at most one formula in Γ', Δ', and Ω'.

Note that $\prec$ is only defined on the extension variables and eigenvariables used so far. Initially, $\prec$ is an ordering where nothing is comparable. As we move down the proof, we order the variables.

The proof is done by induction on the depth of S in the proof π. If we let S be the final sequent, we get a proof of the theorem since $Q_\pi = Q_S$, and conditions 3-5 are the conditions we need for the theorem. Also, note that the induction hypothesis can be stated as a $\Sigma_0^B(PV)$ formula (is a polynomial-time predicate) by saying that the output of the function F on the first i sequents of π meets all of the conditions. This means the induction can be carried out in VPV.

The description of F is done in cases. There is a separate case for each rule of inference. The construction is similar to the proof that extended-Frege p-simulates G_1^* (Theorem 7.48 of [5]). The difference is that the variables need to be ordered.

When the last inference is the $\vee$, $\wedge$, $\neg$ introduction rules, the same rule can be applied in the G_0^* proof we are constructing. It is a simple exercise to check that the induction hypothesis still holds. The same would go for weakening. The other cases are more involved, and are given below.

Inductive Case 1: S is inferred using cut

Suppose

$$S \equiv \Gamma_1, \Gamma_2 \to \Delta_1, \Delta_2, \Omega_1, \Omega_2$$

and that it is derived from

$$S_1 \equiv \exists \vec{x} D(\vec{x}), \Gamma_1 \to \Delta_1, \Omega_1$$

and

$$S_2 \equiv \Gamma_2 \to \Delta_2, \exists \vec{x} D(\vec{x}), \Omega_2,$$

where $D(\vec{x})$ is a Σ_0^q formula. By induction, we have a G_0^* proof of the sequents

$$S_1' \equiv \Lambda_1(\vec{q}), D(\vec{q}), \Gamma_1' \to \Delta_1', \Omega_1'$$

with an ordering $\prec_1$ on $Q_{S_1} \cup E_1$ and

$$S_2' \equiv \Lambda_2, \Gamma_2' \to \Delta_2', D(\vec{e}), \Omega_2'$$

with an ordering $\prec_2$ on $Q_{S_2} \cup E_2$. In this case, we let

$$S' \equiv \Lambda_2, \Lambda_1(\vec{e}), \Gamma'_1, \Gamma'_2 \to \Delta'_1, \Delta'_2, \Omega'_1, \Omega'_2.$$

We say $p_1 \prec p_2$ is true if any of the following conditions hold:

- $p_1 \prec_1 p_2$,
- $p_1 \prec_2 p_2$, or
- $p_1 \in Q_{S_2} \cup E_2$ and $p_2 \in Q_{S_1} \cup E_1$.

We can prove S' by taking the proof of S'_1 and replacing $\vec{q}$ with $\vec{e}$. Because the proofs are treelike, the substitution does not cause any problems. We can then do the cut with S'_1 and S'_2.

We now look at each part of the induction hypothesis to be sure it still holds. Properties 1 to 3 are obvious.

Let us prove property 4. Suppose $e \in E = E_1 \cup E_2$ depends on $p \in Q_S \cup E$. If $e \in E_2$, then e also depends on p in S'_2 since Λ_2 did not change. So, by induction with S_2, $p \prec_2 e$ and, therefore, $p \prec e$. Now suppose $e \in E_1$. If $p \in Q_{S_2} \cup E_2$, then $p \prec e$. If $p \notin Q_{S_2} \cup E_2$, then e depends on p in $\Lambda_1(\vec{q})$ since $\vec{e} \in E_2$. So, by induction with S_1, we get $p \prec_1 e$, which implies $p \prec e$.

Property 5 follows directly form the induction hypothesis. Property 6 follows from the induction hypothesis and the fact that Q_{S_1} and Q_{S_2} are disjoint.

Inductive Case 2: S is inferred using $\forall$-right

Suppose

$$S \equiv \Gamma \to \Delta, \forall y D(y), \Omega$$

and it is derived from

$$S_1 \equiv \Gamma \to \Delta, D(q), \Omega.$$

By induction with S_1, there exists a sequent S'_1 and an ordering $\prec_1$ satisfying the induction hypothesis. In S'_1, there is a formula $D^*(q)$ that is a $(Q_{S_1}, E, \prec_1)$-instance of an $\vee$-expansion of A. That same formula is also a $(Q_S, E, \prec)$-instance of an $\vee$-expansion of $\forall y D(y)$ now that q is part of Q_S.

The construction for this case is fairly simple. We let S' be the same as S'_1, and the ordering $\prec$ is the same as $\prec_1$ except q is now smaller than everything in $Q_{S_1} \cup E$. That is, q is the smallest variable of all variables ordered so far.

Inductive Case 3: S is inferred using $\exists$-right

Suppose

$$S \equiv \Gamma \to \Delta, \exists x D(x), \Omega$$

and it is derived from

$$S_1 \equiv \Gamma, \to \Delta, D(F), \Omega,$$

where F is a Σ_0^q formula.

By induction, we have

$$S'_1 \equiv \Lambda, \Gamma' \to \Delta', D^*(F), \Theta$$

and the ordering $\prec_1$ on $Q_{S_1} \cup E$. Note that we use $D^*(F)$ in place of $D(F)$ because it is possible that $D(F)$ had quantifiers that have already been removed. Since F is a Σ_0^q formula, it would still be intact since (α) does not change Σ_0^q formulas. There are two cases to consider. If $D(F)$ is an ancestor of the final formula then $D^*(F)$ in an instance of $D(F)$. If $D(F)$ is not an ancestor of the final formula, then $D^*(F)$ is $D(F)$ with the existential quantifiers replaced by

extension variables. In either case, $D^*(F)$ is a Σ_0^q formula. The sequent S' will be defined as

$$S' \equiv e \leftrightarrow F, \Lambda, \Gamma' \rightarrow \Delta', D^*(e), \Theta,$$

where e is a new extension variable. As for the ordering $\prec$, it is defined by extending $\prec_1$ by making e the minimum element. Note that $Q_S = Q_{S_1}$.

Let S_2' be

$$e \leftrightarrow F, D'(F) \rightarrow D^*(e).$$

It is easy to derive S_2'. Then it is possible to derive S' from S_1' and S_2' by cutting $D^*(F)$, which is a Σ_0^q formula.

We now look at each part of the induction hypothesis to be sure it still holds. It is easy to see that properties 1 to 3 and property 6 still hold.

For property 4, if e, the new extension variable, depends on p, then p must appear in F, which is part of S_1. This means that $p \notin Q_S \cup E$. So, property 4 holds for e. For other variables, it holds directly from the induction hypothesis.

For property 5, the only instance that changed is $D^*(F)$, assuming $\exists x D(x)$ is not a Σ_1^q formula. Since e replaced the outermost quantifier, e does not have to be larger than any variable, and it is smaller than every variable that replaced inside variables. Therefore $D^*(e)$ is a $(Q_S, E, \prec)$-instance of an $\vee$-expansion of $\exists x D(x)$.

Inductive Case 4: S is inferred using contraction-right

Suppose

$$S \equiv \Gamma \rightarrow \Delta, D, \Omega$$

and it is derived from

$$S_1 \equiv \Gamma \rightarrow \Delta, D, D, \Omega.$$

We look at two different cases: D is a Σ_1^q formula, D is not Σ_1^q.

For the first case, let D be $\exists \vec{x} D'(\vec{x})$, where $D'(\vec{x})$ is a Σ_0^q formula, then, by induction with S_1, we have

$$S_1' \equiv \Lambda, \Gamma' \rightarrow \Delta', D'(\vec{e}^1), D'(\vec{e}^2), \Omega'$$

with the ordering $\prec_1$ on $Q_{S_1} \cup E$. We now have two witnesses for D, and we need to pick the one that works. So, in this case, we let

$$S' \equiv \ldots, e_i^3 \leftrightarrow [(D'(\vec{e}^1) \wedge e_i^1) \vee (\neg D'(\vec{e}^1) \wedge e_i^2)], \ldots, \Lambda, \Gamma' \rightarrow \Delta', D'(\vec{e}^3), \Omega',$$

where $\vec{e}^3$ are new extension variables. The ordering $\prec$ is defined as $\prec_1$ with $\vec{e}^3$ added as the maximum elements so far.

Now we look at each part of the induction hypothesis to be sure it still holds. For properties 1 to 3, notice that the initial part of S' is part of the extension cedent and it defines the new variables $\vec{e}^3$. With this observation, it is easy to see that properties 1 to 3 still hold. Now to look at property 4. Since $\vec{e}^3$ are the largest elements in the ordering $\prec$, anything they depend on must be incomparable or smaller. So property 4 holds for $\vec{e}^3$. For other extension variables, it holds directly from the induction hypothesis. Property 5 follows directly from the induction hypothesis since Ω did not change. Property 6 follows directly from the induction hypothesis.

For the second possibility, assume that D is not a Σ_1^q formula. The by induction with S_1 we get

$$S_1' \equiv \Lambda, \Gamma' \to \Delta', D_1^*, D_2^*, \Omega'$$

where D_1^* and D_2^* are $(Q_{S_1}, E, \prec_1)$-instances of an $\vee$-expansion of D. Then we let

$$S' \equiv \Lambda, \Gamma' \to \Delta', D_1^* \vee D_2^*, \Omega'$$

which can be obtained from S_1' using $\vee$-right. Notice that $D_1^* \vee D_2^*$ is also a $(Q_S, E, \prec)$-instance of an $\vee$-expansion of D, and, since the ordering is not changed, the induction hypothesis still holds.

Inductive Case 5: S is inferred using contraction-left

Suppose

$$S \equiv D, \Gamma \to \Delta, \Omega$$

and it is derived from

$$S_1 \equiv D, D, \Gamma, \to \Delta, \Omega.$$

The formula D must be a Σ_1^q formula and an ancestor of a cut formula. Let it be of the form $\exists x F(x)$, where $F(x)$ is a Σ_0^q formula. For now, we assume that D has a single existential quantifier, but the construction easily generalizes. By induction with S_1, there exists a sequent S_1' of the form

$$\Lambda(q_1, q_2), F(q_1), F(q_2), \Gamma' \to \Delta', \Omega'$$

with an ordering $\prec$. Without lose of generality, assume that $q_1 \prec q_2$. Then, we let

$$S' \equiv \Lambda(q_1, q_1), F(q_1), \Gamma' \to \Delta', \Omega'$$

The ordering will remain the same.

To prove S', we take the proof of S_1', replace every instance of q_2 by q_1, and then contract the two copies of $F(q_1)$. The substitution can be done because the proof is treelike.

We now look at each part of the induction hypothesis to be sure it still holds. It is easy to see that properties 1 to 3 hold.

Property 4 follows from the induction hypothesis. Note that, if a variable depended on q_2, it now depends on q_1, but that is fine since q_1 is smaller.

Property 5 holds since Ω did not change. ⊣

From this, we are able to prove the witnessing theorem (Theorem 3.3). We will restate it here.

THEOREM 3.11 (KPT Witnessing for G_1^*). *There exists a PV (polynomial-time) function F such that VPV proves the following. Let π be a G_1^* proof of a prenex Σ_2^q formula $A(\vec{p}) \equiv \exists \vec{x} \forall \vec{y} B(\vec{x}, \vec{y}, \vec{p})$, where $B(\vec{x}, \vec{y}, \vec{p})$ is a Σ_0^q formula with all free variables shown. Then, given π, F outputs a G_0^* proof of a sequent $\Lambda \to \Theta$ where*

1. *Θ is a series of formulas of the form $B(\vec{e}^i, \vec{q}^i, \vec{p})$, where $\vec{e}^i \in E$*
2. *Λ is an extension cedent defining a new set of variables E in terms of $\vec{q}^1, \ldots, \vec{q}^n$ and $\vec{p}$,*
3. *$\vec{e}^i$ does not depend on $\vec{q}^j$, for $j \geq i$, and*
4. *$\vec{q}^i$ and $\vec{q}^j$ are disjoint.*

PROOF. By the Herbrand theorem above, there is a proof of $\Lambda \to A^*$, where A^* is a $(Q_\pi, E, \prec)$-instance of an $\vee$-expansion of A. We need to show how to get Θ from A^*.

The first observation we make is that A^* is of the from

$$B(\vec{e}^1, \vec{q}^1, \vec{p}) \vee \ldots \vee B(\vec{e}^n, \vec{q}^n, \vec{p}).$$

This is true because the rule (α) gives multiple copies of B, which all remain in tact since B is a Σ_0^q formula, combined using $\vee$. This means that A^* is essentially the Θ we want.

Without lose of generality, we can assume that, if $i < j$, then the smallest variable in $\vec{q}^i$ is smaller than the smallest variable in $\vec{q}^j$. This implies that $\vec{e}^i$ does not depend on $\vec{q}^j$ for $j \geq i$. This is because every variable in $\vec{e}^i$ is smaller than every variable in $\vec{q}^i$ since we have a $(Q_\pi, E, \prec)$-instance.

The final observation is that if $\vec{q}^i$ and $\vec{q}^j$ contain a common variable then $\vec{e}^i$ and $\vec{e}^j$ must be the same. Otherwise, if $\vec{e}^i$ and $\vec{e}^j$ are different, then an application of (α) must have occurred that would make part of $\vec{e}^i$ and $\vec{e}^j$ correspond to different existential variables. Since the universal variables are in the scope of these existential variables, $\vec{q}^i$ and $\vec{q}^j$ would correspond to different quantifiers making them disjoint.

If we have that $\vec{q}^i$ and $\vec{q}^j$ are not disjoint, we are able to replace $\vec{q}^j$ by $\vec{q}^i$ everywhere in the proof. Then we can contract the two copies of $B(\vec{e}^i, \vec{q}^i)$. $\dashv$

3.2. Witnessing for G_i^*. In the statement of the original KPT witnessing theorem for VPV, polynomial-time functions are used to find the possible witnesses; however, for TV^i, the KPT witnessing theorem uses functions in $FP^{\Sigma_i^p}$. Corresponding to that we will generalize the definition of an extension cedent to allow oracles from the other levels of the polynomial hierarchy.

DEFINITION 3.12 (i-extension cedent). *An* i-extension cedent *is a series of formulas of the form*

$$e_1 \leftrightarrow E_1, e_2 \leftrightarrow E_2, \ldots, e_n \leftrightarrow E_n$$

such that E_m *is a* $\Sigma_i^q \cup \Pi_i^q$ *formula that does not mention the variables* $e_m, e_{m+1}, \ldots, e_n$.

Note that an extension cedent is the same as a 0-extension cedent.

DEFINITION 3.13 (i-expansion). *The same as an* $\vee$*-expansion except that* B *must be a non-*$(\Sigma_i^q \cup \Pi_i^q)$ *formula instead of a non-*Σ_0^q *formula.*

DEFINITION 3.14 ($(i, Q, E, \prec)$-instance). *An* $(i, Q, E, \prec)$-instance *of a formula* A *is the same as a* $(Q, E, \prec)$*-instance of* A *except that* $\Sigma_i^q \cup \Pi_i^q$ *subformulas of* A *are not changed. That is, the quantifiers that appear in the* i *innermost blocks of quantifiers are not replaced.*

Note that a $(i, Q, E, \prec)$-instance of a formula will always be a $\Sigma_0^q(\Sigma_i^q)$ formula.

EXAMPLE 3.15. *Let* A *be the formula* $\exists\vec{x}\forall\vec{y}\exists\vec{z}B(\vec{x}, \vec{y}, \vec{z})$*, where* B *is a* Σ_0^q *formula. Then a* $(1, Q, E, \prec)$*-instance of* A *would be* $\exists\vec{z}B(\vec{e}, \vec{q}, \vec{z})$ *and a* $(2, Q, E, \prec)$*-instance would be* $\forall\vec{y}\exists\vec{z}B(\vec{e}, \vec{y}, \vec{z})$.

THEOREM 3.16 (Herbrand Theorem for G_i^*). *For each $i > 0$, there is a PV function F such that VPV proves the following. Let π be a G_i^* proof of A. Then, given π, F outputs a G_{i-1}^* proof of a sequent $\Lambda \to A^*$ and a total ordering $\prec$ of the variables $Q_\pi \cup E$, where E is a set of variables that do not appear in π, with the following properties:*

- *Λ is a $(i-1)$-extension cedent defining the variables in E in terms of Q_π and the free variables of A;*
- *for $e \in E$, if e depends on a variable $p \in Q_\pi \cup E$, then $p \prec e$; and*
- *A^* is an $(i-1, Q_\pi, E, \prec)$-instance of an $(i-1)$-expansion of A.*

Before we prove this theorem, we should note that it does not seem like we can improve the complexity of the extension cedent. For, if the $(i-1)$-extension cedent could be replaced by a $(i-2)$-extension cedent, this could be used to show that G_{i-1}^* p-simulates G_i^* for prenex formulas. The proof would be similar to the proof of Theorem 4.2.

PROOF OF THEOREM 3.16. The proof is almost the same as in the G_1^* case. The quantifier complexity of the cut formulas is reduced by witnessing the outermost block of existential quantifiers with extension variables. The only difference is that we can no longer skip all of the quantifier introduction rules. Some will have to be added. For example, if we have a cut formula $\exists\vec{x}\forall\vec{y}C(\vec{x},\vec{y})$, then we will replace $\vec{x}$ by extension variables, but we still add $\forall\vec{y}$ to the formula.

This construction can be described more formally. As before, each sequent S in π can be divided into three parts: Γ which contains all of the formulas on the left-hand side; Δ, which contains the formulas on the right-hand side that are ancestors of cut formulas; and Ω, which contains the ancestors of the final formula on the right-hand side. Note that by the subformula property, we know that Γ and Δ contain only Σ_i^q formulas. For each sequent $S \equiv \Gamma \to \Delta, \Omega$ in π, we construct a G_{i-1}^* proof of a sequent

$$S' \equiv \Lambda, \Gamma' \to \Delta', \Omega',$$

and a total ordering $\prec$ on $Q_S \cup E$ where

1. Γ' is obtained from Γ by replacing each non-Σ_{i-1}^q formula $\exists\vec{z}D(\vec{z})$ by $D(\vec{q})$, where D is Π_{i-1}^q and $\vec{q} \in Q_S$. (We use different $\vec{q}$ for different formulas.)
2. Δ' is obtained from Δ be replacing each non-Σ_{i-1}^q formula $\exists\vec{z}D(\vec{z})$ by $D(\vec{e})$, where D is Π_{i-1}^q and $\vec{e} \in E$. (We use different $\vec{e}$ for different formulas.)
3. Λ is an $(i-1)$-extension cedent defining E;
4. for $e \in E$, if e depends on a variable $p \in Q_S \cup E$, then $p \prec e$;
5. Ω' is obtained from Ω by replacing each formula B by an $(i-1, Q_S, E, \prec)$-instance of an $(i-1)$-expansion of B; and
6. each $q \in Q_S$ appears in at most 1 formula in Γ', Δ', and Ω'.

The construction is the same as in Theorem 3.10 except for the need to add a few new cases. If $\exists$-right is applied with a Σ_{i-1}^q principal formula or $\forall$-left is applied with a Π_{i-1}^q principal formula, the same inference can be used in the G_{i-1}^* proof we are constructing. We must also consider when S is derived using $\exists$-left with a Σ_{i-1}^q principal formula and when S is derived using $\forall$-right with a Π_{i-1}^q principal formula. Both cases are handled in the same way, so we only describe the latter.

Suppose

$$S \equiv \Gamma \to \Delta, \forall q D(q), \Omega$$

and it is derived from

$$S_1 \equiv \Gamma \to \Delta, D(q), \Omega$$

where $D(q)$ is Π^q_{i-1}. By induction, we have a G^*_{i-1} proof of S'_1, where

$$S'_1 \equiv \Lambda(q), \Gamma' \to \Delta', \Omega', D(q).$$

We know q does not appear in Γ' or Δ' since is was used as an eigenvariable, but it is still possible that the extension variables depend on it, in which case it would appear in Λ.

The first step is to replace q by a new extension variable e. This gives

$$\Lambda(e), \Gamma' \to \Delta', \Omega', D(e). \tag{3.1}$$

We then derive

$$e \leftrightarrow D(\bot), D(e) \to \forall q D(q). \tag{3.2}$$

See Lemma 3.17 below for how to do this. We finish by deriving

$$e \leftrightarrow D(\bot), \Lambda(e), \Gamma' \to \Delta', \Omega', \forall q D(q)$$

by cut with sequents (3.1) and (3.2) and cut formula $D(e)$.

The ordering is changed by adding e as the smallest element. Note that, since $D(q)$ is in S_1, $D(\bot)$ does not contain any extension variables or eigenvariables in Q_S. With this fact in mind, we can see all of the conditions in the induction hypothesis follow. ⊣

LEMMA 3.17. *Let $B(q)$ be a Σ^q_i or Π^q_i formula. Then there exists polynomial-size G^*_i proofs of the sequents*

$$e \leftrightarrow B(\bot), B(e) \to \forall q B(q)$$
$$e \leftrightarrow B(\top), \exists q B(q) \to B(e)$$

PROOF. The proof for the two sequents are essentially the same, so we only give the construction for the first one. Informally, the reason the first sequent is true is that we are picking a value for e that makes $B(e)$ false if possible. So, if $B(\bot)$ is false, we make e false, otherwise we make e true, which is the only other possible value.

First, it is possible to get cut-free proofs of the following four sequents.

$$e, B(e) \to B(\top)$$
$$B(e) \to B(\bot), e$$
$$e, B(\top) \to B(e)$$
$$B(\bot) \to B(e), e$$

This can be shown by simultaneous structural induction on the formula $B(e)$.

We use this in the following derivation:

$$\text{Cut } q\ \frac{q, B(\top) \to B(q) \qquad B(\bot) \to B(q), q}{\forall\text{-right}\ \dfrac{B(\bot), B(\top) \to B(q)}{B(\bot), B(\top) \to \forall q B(q)}}$$

So to finish proving this lemma, all we need are proofs of the sequents

$$e \leftrightarrow B(\bot), B(e) \to B(\bot)$$
$$e \leftrightarrow B(\bot), B(e) \to B(\top)$$

We prove the first one as follows:

$$\text{Weakening and } \vee\text{-left } \frac{\text{Cut e } \dfrac{e, e \supset B(\bot) \to B(\bot) \qquad B(e) \to B(\bot), e}{e \supset B(\bot), B(e) \to B(\bot)}}{e \leftrightarrow B(\bot), B(e) \to B(\bot)}$$

The second one is proved as follows:

$$\text{Cut } B(\bot)\ \frac{B(e) \to e, B(\bot) \qquad B(\bot) \supset e, B(\bot) \to e}{\text{Cut } e\ \dfrac{B(\bot) \supset e, B(e) \to e \qquad\qquad e, B(e) \to B(\top)}{B(\bot) \supset e, B(e) \to B(\top)}}$$

Note that the only cut formulas are e, $B(\bot)$, and $B(\top)$; therefore, the proof is a G_i^* proof. $\dashv$

§4. GPV^* **and** G_1^***.** We now move on to applications of the Herbrand Theorem for G_1^*. The first application deals with a seemingly weaker proof system.

DEFINITION 4.1. *The proof system GPV^* is G_1^* in which cut formulas are restricted to Σ_0^q formulas or formulas of the form $\exists x[x \leftrightarrow A]$, where A is a Σ_0^q formula that does not mention x.*

At first glance, it seems like GPV^* would be a weaker proof system than G_1^* because the cut formulas are less expressive. The cut formulas in GPV^* can be trivially witnessed, but the cut formulas in G_1^* are NP-hard. Nevertheless, it can be shown that GPV^* and G_1^* are p-equivalent for prenex formulas. One direction is easy since every GPV^* proof is a G_1^* proof, so all that is left is to prove the other direction.

THEOREM 4.2. *VPV proves that GPV^* p-simulates G_1^* for prenex formulas.*

PROOF. Let π be a G_1^* proof of a formula A of the form

$$\forall \vec{y}^0 \exists \vec{x}^1 \forall \vec{y}^1 \ldots \exists \vec{x}^n \forall \vec{y}^n B(\vec{y}^0, \vec{x}^1, \vec{y}^1, \ldots, \vec{x}^n, \vec{y}^n),$$

where B is a Σ_0^q formula. By the Theorem 3.10, VPV proves that there exists a G_0^* proof π' of a sequent $\Lambda \to A^*$ and a total ordering $\prec$ of the variables $Q_\pi \cup E$ meeting the conditions of the theorem. Since A is in prenex form, we know that A^* is of the form $\bigvee_{i=0}^m B(\vec{q}^{i,0}, \vec{e}^{i,1}, \vec{q}^{i,2}, \ldots, \vec{e}^{i,n}, \vec{q}^{i,n})$. From this we are able to get a proof of $\Lambda \to \Theta$ where

$$\Theta \equiv B(\vec{q}^{0,0}, \vec{e}^{0,1}, \vec{q}^{0,2}, \ldots, \vec{e}^{0,n}, \vec{q}^{0,n}), \ldots, B(\vec{q}^{m,0}, \vec{e}^{m,1}, \vec{q}^{m,2}, \ldots, \vec{e}^{m,n}, \vec{q}^{m,n})$$

by deriving $A^* \to \Theta$ and cutting A^*.

We describe an algorithm that takes as input π' and $\prec$. The algorithm extends π' into a GPV^* proof of A. At any stage, π' will be a proof of a sequent $\Lambda' \to \Theta'$, where Λ' is a subsequence of Λ and Θ' is a sub-series of Θ with some quantifiers added. The algorithm has four steps:

Step 1: Add as many existential quantifiers to the formulas in Θ' as possible using $\exists$-right rules such that the formula is still a subformula of A.
Step 2: Use contraction to combine as many formulas in Θ' as possible.
Step 3: Find the largest variable that is mentioned in Λ' or Θ'.
- Step 3a: If it is an extension variable e, apply $\exists$-left to the formula $e \leftrightarrow E$ with e as the eigenvariable. Then cut the formula $\exists e[e \leftrightarrow E]$ after deriving $\rightarrow \exists e[e \leftrightarrow E]$.
- Step 3b: If it was an eigenvariable q in π, then apply $\forall$-right with q as the eigenvariable.

Step 4: Repeat steps 1 to 3 until there is no change.

At first, it may not be obvious that this algorithm works. For example, it is not obvious that the eigenvariable restriction for $\exists$-left or $\forall$-right rules in Step 3 is met. To show that the eigenvariable restriction is met, we make two observations. First, if p is the largest variable in Λ' and Θ', then no extension variable can depend on p. Otherwise, that variable would be larger than p. Second, if we are in Step 3 and p is the largest variable in Λ' and Θ', then p cannot be mentioned in Θ' unless it is in Q_π; otherwise p would be an extension variable and have been used as the target formula in an $\exists$-right rule in Step 1. If this is not the case, an eigenvariable that appears to the right of p is still present, and this variable must be larger than p. For the same reason, we know that there cannot be two formulas in Θ with p replacing a universal variable that have not been contracted yet. This means the eigenvariable restriction is met in Step 3.

When the algorithm is done, we will have a proof of the formula we want. Notice that Λ' would be empty because every extension variable has been removed. Also, Θ' would be the single formula A since every formula in Θ would have every quantifier added by now, and would have been contracted to a single formula. We know the algorithm eventually stops because we continually reduce the number of variables in π'. ⊣

§5. G_i and G^*_{i+1}. As has already been mentioned, for $i > 0$, G_i is commonly connected with the theory TV^i and G^*_{i+1} is commonly connected with V^{i+1}. Since the two theories have the same Σ^B_{i+1} theorems, it was natural that the two proof systems are p-equivalent when proving Σ^q_{i+1} formulas. However, we want to extend this to more general formulas. In [13], it was shown that one direction is probably not possible. Namely that, under an appropriate complexity assumption, there is a family of Σ^q_{i+2} formulas for which G^*_{i+1} does not p-simulate G_i. Here we prove that G_i p-simulates G^*_{i+1} for all formulas.

The proof is based on the proof of Krajicek that depth d, DAG-like PK can p-simulate depth $d+1$, treelike PK. The observation of the similarity between the two theorems is due to Toni Pitassi.

Definition 5.1 (The i-Substitution Rule). *The i-substitution rule is*

$$\frac{A_1(p), \ldots, A_m(p), \Gamma \to \Delta, B_1(p), \ldots, B_n(p)}{A_1(C), \ldots, A_m(C), \Gamma \to \Delta, B_1(C), \ldots, B_n(C)}$$

where C is a quantifier-free formula, $A_1, \ldots, A_m, B_1, \ldots, B_n$ are $\Sigma_i^q \cup \Pi_i^q$ formulas, and p does not appear in the bottom sequent.

LEMMA 5.2. *For $i > 0$, G_i^* p-simulates the i-substitution rule.*

PROOF. The proof is the same as the proof of Lemma 2.1 in [9]. We will describe how to do the simulation for the case where there is one A and B. The general case is done the same way.

Suppose we have a derivation of

$$A(p), \Gamma \to \Delta, B(p). \tag{5.1}$$

We want to derive

$$A(C), \Gamma \to \Delta, B(C).$$

First we derive

$$p \leftrightarrow C, A(C) \to A(p),$$

and cut this with (5.1), where $A(p)$ is the cut formula. This gives

$$p \leftrightarrow C, A(C), \Gamma \to \Delta, B(p). \tag{5.2}$$

Then we derive

$$p \leftrightarrow C, B(p) \to B(C),$$

and cut this with (5.2), where $B(p)$ is the cut formula. This gives

$$p \leftrightarrow C, A(C), \Gamma \to \Delta, B(C). \tag{5.3}$$

We then apply $\exists$-left to this sequent with p as the eigenvariable, and then cut $\exists p[p \leftrightarrow C]$ after deriving $\to \exists p[p \leftrightarrow C]$. ⊣

THEOREM 5.3. *For $i > 0$, G_i p-simulates G_{i+1}^*.*

PROOF. At a high level, this proof is done by carefully applying one step of cut-elimination to each cut formula. The increase in the size of the proof in the cut-elimination theorem comes from repeating part of the proof multiple times. We avoid this increase by creating a DAG-like proof.

Let π be a G_{i+1}^* proof. The reason π is not a G_i proof is that it would contain cut formulas that are not Σ_i^q or Π_i^q. We can assume these formulas are Σ_{i+1}^q and are of the form

$$\exists x_1 \ldots \exists x_n C(x_1, \ldots, x_n),$$

where C is Π_i^q. We can assume this because, in [12], Morioka proved that we all G_{i+1}^* proofs can be transformed into a G_{i+1}^* proof where the cut formulas are prenex. We need to turn these cut formulas into Π_i^q cut formulas. To do this, we change all of the non-$(\Sigma_i^q \cup \Pi_i^q)$ formulas that are ancestors of these cut formulas. These formulas are of the form

$$\exists x_l \ldots \exists x_n C(D_1, \ldots, D_{l-1}, x_l, \ldots, x_n), \tag{5.4}$$

where D_j is a Σ_0^q formula for $j < l$, and $C(\vec{x})$ is a Π_i^q formula. Note that, if this formula is on the left side of a sequent, then the formula D_i will actually be variables that eventually get used as eigenvariables in an $\exists$-left rule. From now

on, we will assume all formulas of the form (5.4) are ancestors of cut formulas. Those that are not are simply ignored.

The construction will be done inductively. We start with the first sequent in π and work our way down the proof. For each sequent $S \equiv \Gamma \to \Delta$ in π, we give a G_i proof π' of a sequent $S' \equiv \Gamma' \to \Delta'$ where

1. Γ' is obtain from Γ by replacing every formula of the form (5.4) by $C(D_1, \ldots, D_{l-1}, x_l^C, \ldots, x_n^C)$,
2. Δ' is obtained from Δ by removing every formula of the form (5.4),
3. the sequent
$$C(D_1, \ldots, D_{l-1}, x_l^C, \ldots, x_n^C) \to$$
can be used as an axiom if and only if Δ contains a formula of the form (5.4).

For example, if S is the sequent

$$\exists x_2, x_3 C_1(q_1, x_2, x_3), \Gamma \to \Delta, \exists x_3, x_4 C_2(D_1, D_2, x_3, x_4),$$

S' would be

$$C_1(q_1, x_2^{C_1}, x_3^{C_1}), \Gamma \to \Delta,$$

and when we prove S', we are allowed to use

$$C_2(D_1, D_2, x_3^{C_2}, x_4^{C_2}) \to$$

as an axiom. In essence, we are saying, if we can derive

$$C_2(D_1, D_2, x_3^{C_2}, x_4^{C_2}) \to,$$

we can prove S'. Note that, when we get to the final sequent, no formula is an ancestor of a cut formula. Therefore, if S is the final formula in π, $S' = S$ and the only initial sequents are of the form $x \to x$. So this will give us a proof of the theorem.

The construction of π' is given inductively. There is a separate case for each rule of inference. Most cases are simple and are left to the reader. The only cases we will give are cut, $\exists$-left, and $\exists$-right.

Cut: Suppose $S \equiv \Gamma \to \Delta$ is derived from S_1 and S_2 using cut. Let the cut formula be $\exists \vec{x} C(\vec{x})$. By induction with S_1, we have a G_i proof π_1' of

$$S_1' \equiv C(\vec{x}^C), \Gamma' \to \Delta'.$$

By induction with S_2, we have a G_i proof π_2' of $\Gamma' \to \Delta'$ using the axiom $C(\vec{x}^C) \to$. Notice that π_2' is a proof of the sequent we want, but it uses an axiom we are no longer able to use. However, π_1' gives us a derivation of this axiom, with a few extra formulas.

The first step in the construction of π' is to add Γ' to the left and Δ' to the right of every sequent in π_2'. This makes the axiom we want to remove $C_i(\vec{x}^i), \Gamma' \to \Delta'$, which is the final sequent π_1'. So, π' is π_1' followed by the new π_2'. Note that the axiom would have been used once for every time $\exists x_n$ was introduced in the original proof. Each of these formulas would later be contracted into the single cut formula. However, since we are constructing a DAG-like proof, we do not need to repeat π_1' multiple times. This gives a proof of $\Gamma', \Gamma' \to \Delta', \Delta'$, from which we can derive $\Gamma' \to \Delta'$ using contraction.

∃-left: Suppose S is

$$\exists x_j \ldots \exists x_n C(q_1, \ldots, q_{j-1}, \mathbf{x_j}, x_{j+1}, \ldots, x_n), \Gamma \to \Delta,$$

and it was derived from S_1

$$\exists x_{j+1} \ldots \exists x_n C(q_1, \ldots, q_{j-1}, \mathbf{q_j}, x_{j+1}, \ldots, x_n), \Gamma \to \Delta.$$

By induction with S_1, we get a G_i proof of

$$C(q_1, \ldots, q_{j-1}, \mathbf{q_j}, x_{j+1}^C, \ldots, x_n^C), \Gamma' \to \Delta'.$$

Since q_j was used as an eigenvariable, it only appears in that one formula. Therefore we can replace q_j by x_j^C using the i-substitution rule. This gives us π'.

∃-right: Suppose S is

$$\Gamma \to \Delta, \exists x_j \ldots \exists x_n C(D_1, \ldots, D_{j-1}, \mathbf{x_j}, x_{j+1}, \ldots, x_n),$$

and it was derived from S_1

$$\Gamma \to \Delta, \exists x_{j+1} \ldots \exists x_n C(D_1, \ldots, D_{j-1}, \mathbf{D_j}, x_{j+1}, \ldots, x_n).$$

First assume $j < n$. That is we had at least one quantifier already. By induction with S_1, we get a G_i proof of $\Gamma' \to \Delta'$ using the axiom

$$C(\ldots, \mathbf{D_j}, \ldots) \to . \tag{5.5}$$

We cannot use this axiom anymore. Instead, we use the axiom

$$C(\ldots, \mathbf{x_j^C}, \ldots) \to$$

and derive (5.5) using the i-substitution rule.

If $j = n$, the construction is a little different. By induction with S_1, we get a G_i proof of

$$\Gamma' \to \Delta', C(\ldots, D_{n-1}, D_n). \tag{5.6}$$

To construct π', we take the axiom we can now use,

$$C(\ldots, D_{n-1}, x_n^C) \to,$$

and derive

$$C(\ldots, D_{n-1}, D_n) \to$$

using the i-substitution rule. Then we cut with (5.6). ⊣

§6. Reflection Principles. We can also use the Herbrand Theorem to prove reflection principles. Proving reflection principles is the standard method of assessing the strength of a proof system relative to a theory. For example, the Σ_1^q reasoning of G_1^* is not stronger than the Σ_1^B reasoning of V^1 because V^1 proves Σ_1^q-RFN(G_1^*) [7]. Our goal is to find the weakest fragment of V that proves Σ_i^q-RFN(G_1^*). In [12], it was shown that TV^0 does not prove Σ_2^q-RFN(G_1^*) unless the polynomial-time hierarchy collapses. Using the same ideas, it is possible to show that TV^i does not prove Σ_{i+2}^q-RFN(G_1^*), for $i \geq 0$, unless the polynomial-time hierarchy collapses. This still leaves open whether or not V^i proves Σ_{i+1}^q-RFN(G_1^*) for $i \geq 1$. We prove that, in fact, it does.

We first prove the simplest case. Namely, that V^1 proves (prenex Σ_2^q)-RFN(G_1^*). The proof serves as a template for the general case, which we prove right after.

THEOREM 6.1. *V^1 proves (prenex Σ^q_2)-RFN(G^*_1).*

PROOF. Let π be a G^*_1 proof of a prenex Σ^q_2 formula A. So A is of the form

$$\exists \vec{x} \forall \vec{y} B(\vec{x}, \vec{y}, \vec{p}),$$

where B is a Σ^q_0 formula. In this formula, $\vec{p}$ is all of the free variables in A, and should be understood as being implicitly universally quantified. We want to prove in V^1 that, given values for $\vec{p}$, there exists values for $\vec{x}$ that witness the formula.

By the KPT witnessing theorem for G^*_1 (Theorem 3.3), V^1 proves that there is a G^*_0 proof of a sequent

$$S \equiv \Lambda \to \Theta,$$

meeting the conditions of the theorem.

Let

$$\begin{aligned}
&\psi(m, \Lambda, \Theta, P) \equiv \\
&\exists E\ \exists Q \text{ “E is a truth assignment to the extension variables”} \\
&\quad \wedge \text{ “Q is a truth assignment to the eigenvariables”} \\
&\quad \wedge \forall i < m\ (P \cup E \cup Q) \models_0 \neg B(\vec{e}^{\,i}, \vec{q}^{\,i}, \vec{p}) \\
&\quad \wedge\ (P \cup E \cup Q) \models_0 \Lambda
\end{aligned}$$

This formula says that there exists assignments E and Q that satisfy Λ and make the first m formulas in Θ false. It is easy to bound the size of E and Q. This means that ψ is equivalent to a Σ^B_1 formula.

Using Σ^B_1-MAX, we find the maximum value m_0 for m that satisfies ψ given values for Λ, Θ, and P. Then $\vec{e}^{\,m_0+1}$ are the witnesses we are looking for, which we now prove.

First note that $\psi(0)$ is true. We can set Q to the assignment that sets everything to false, and compute E that satisfies Λ. Also note that $m_0 < n$ since it is not possible the falsify all of the formulas in Θ. This would violate the Σ^q_0-RFN(G^*_0), which is provable in V^1. This means that $\vec{e}^{\,m_0+1}$ exists.

Let E and Q be witnesses for $\psi(m_0)$. For the sake of contradiction assume $\vec{e}^{\,m_0+1}$ is not a witness for $\exists \vec{x} \forall y B(\vec{x}, \vec{y}, \vec{p})$. Change Q so that $\vec{q}^{\,m_0+1}$ are assigned values falsifying $B(\vec{e}^{\,m_0+1}, \vec{q}^{\,m_0+1}, \vec{p})$. We can then change E so that Λ is satisfied. Since $\vec{e}^{\,j}$, for $j \leq m_0 + 1$, does not depend on $\vec{q}^{\,m_0+1}$, their values stay the same. This means we now have E and Q making the first $m_0 + 1$ formula in Θ false, violating our choice of m_0. ⊣

THEOREM 6.2. *$V^i \vdash \Sigma^q_{i+1}$-RFN($G^*_i$).*

PROOF. Suppose we have a G^*_i proof of a Σ^q_{i+1} formula A. By Theorem 3.16, we can find a G^*_{i-1} proof of an instance of an $\vee$-expansion of A, with and ordering $\prec$. Let A^* be the $\vee$-expansion of A. Then, by Lemma 6.3 below, all we need to do is prove A^* is valid in order to prove the reflection principle.

This is done in a similar fashion to Theorem 6.1. In the previous case, a block of quantifiers was $\vec{q}^{\,i}$. In this theorem, we put all of the universal quantifiers that are in the scope of the same existential quantifiers in one group. For example, if

$$A^* \equiv (\exists x_1 (\forall y_1 B \vee \forall y_2 B)) \wedge \exists x_2 (\forall y_3 C \wedge \exists x_3 \forall y_4 D)$$

then there are three groups of quantifiers. The variables y_1 and y_2 form one group since they are both in the scope of x_1 and no other variables. The variable y_3 forms the second group. It cannot be in the same group as y_4 because it is not in the scope of x_3. The final group is y_4.

We order the groups of universal variables by the smallest variable in the group, and for each group we associate the formula where the variables for the smaller groups have been replaced. For example, we use A^* above and the instance is

$$A' \equiv (B(e_1, q_1) \vee B(e_1, q_2)) \wedge (C(e_2, q_3) \wedge D(e_2, e_3, q_4))$$

where $q_3 \prec q_1 \prec q_4 \prec q_2$. Then the formula associated with the group $\{q_1, q_2\}$ is

$$(B(e_1, q_1) \vee B(e_1, q_2)) \wedge (C(e_2, q_3) \wedge \exists x_3 \forall y_4 D(e_2, x_3, y_4))$$

We removed the quantifiers $\forall y_3$, $\forall y_2$, and $\forall y_1$, plus all existential quantifiers that are outside (smaller than) these universal quantifiers. Note we did not replace y_4 with q_4 because the group $\{q_4\}$ is larger then the group $\{q_1, q_2\}$, even if q_2 is larger than q_4.

Then, by Σ_i^B-MAX, we are able to find values for the eigenvariables that make as many of these formulas false, starting with the formula for the first group and going through the groups in order. By Σ_{i-1}^q-RFN(G_{i-1}^*), which is provable in V^i, we cannot make the last formula false. As before we are able to extract the witness we want. ⊣

LEMMA 6.3. *V^i proves that, if A^* is an $\vee$-expansion of a Σ_i^q formula A, then*

$$\sigma \models_i A^* \leftrightarrow \sigma \models_i A.$$

PROOF. Done by induction on the number of applications of (α) (Definition 3.5) used to obtain A^* from A. ⊣

§7. New Axiomatization of V. In this section, we will strengthen a result from [9]. In that paper, Krajicek and Pudlak showed that V can be axiomatized by $V^1 + \{\Sigma_i^q\text{-RFN}(G_i) \mid i \in \mathbb{N}\}$. A similar proof can be used to prove that V can be axiomatized by $V^1 + \{\Sigma_i^q\text{-RFN}(G_i^*) \mid i \in \mathbb{N}\}$. In this section, we show that V can also be axiomatized by $V^1 + \{\Sigma_i^q\text{-RFN}(CFG^*) \mid i \in \mathbb{N}\}$, where CFG^* is the cut-free version of G^*. Note that CFG^* is a weaker proof system than any of the other fragments of G.

Just a bit of notation. If A is a formula with free variables $\vec{p}$, then $\exists A$, called the existential closure of A, is the formula $\exists \vec{p} A$.

LEMMA 7.1. *V^1 proves*

$$\Sigma_{i+1}^q\text{-}RFN(CFG^*) \leftrightarrow \Sigma_{i+1}^q\text{-}RFN(G_i^*).$$

PROOF. The if direction is easy since a CFG^* proof is also a G_i^* proof. The only if direction is not as easy. Assume Σ_{i+1}^q-RFN(CFG^*), and argue in V^1. Given a G_i^* proof π of a Σ_{i+1}^q formula A, we change it into a CFG^* proof of a formula

$$B \equiv A \vee \bigvee_{j=1}^{n} \exists (C_j \wedge \neg C_j),$$

where $C_1, \ldots, C_n$ are all of the cut formulas in π.

This is done by first replacing each cut by

$$\dfrac{\dfrac{\Gamma \to \Delta, C \qquad \dfrac{C, \Gamma \to \Delta}{\Gamma \to \Delta, \neg C}}{\Gamma \to \Delta, C \wedge \neg C}}{\Gamma \to \Delta, \exists(C \wedge \neg C)}$$

The sequents in the rest of the proof are changed to include $\exists(C_i \wedge \neg C_i)$. Note that none of the inferences are affected by adding this formula. The only problem could be the eigenvariable restriction in $\exists$-left and $\forall$-right inferences; however, since the new formula does not have any free variables, there is no problem. At the end of the proof, the A is combined with the new formulas using $\vee$-right inferences.

Since the cut formulas are Σ^q_i formulas, B is a Σ^q_{i+1} formula. By Σ^q_{i+1}-RFN(CFG^*), B is true, and, since $\exists(C_i \wedge \neg C_i)$ cannot be true, A must be true. This can be done in V^1 since it proves the Tarski conditions for the true definition. ⊣

COROLLARY 7.2. $V = V^1 + \{\Sigma^q_i\text{-}RFN(CFG^*) \mid i \in \mathbb{N}\}$.

PROOF. Follows from the lemma above, Krajicek and Pudlak's axiomatization of V, and the fact that Σ^q_{i+1}-RFN(G^*_i) implies Σ^q_i-RFN(G^*_i) ⊣

§8. Conclusion And Future Work. In this paper, we looked at G_i and G^*_i as proof systems for proving formulas with high quantifier complexity. Many questions of this type have been resolved, and, as a whole, these results indicate that power of these proof systems grows beyond any finite level of the V (or S_2) hierarchy.

The Herbrand theorem for G^*_i provides an interesting reduction of the provability of high complexity formulas to the provability of a Σ^q_{i+1} formula. Only a couple corollaries have a been given, but we believe it is worth further exploring the use of this theorem as a tool in the proof theory of quantified propositional proof systems.

We still need to explore these proof systems as proof systems for low quantifier complexity formulas. One interesting problem would be the find the complexity of the witnessing problem for these proof systems. For example, we could ask how hard it is to find a witness for Σ^q_1 formula given a G^*_i proof of the formula. It can be shown that this problem is equivalent to the corresponding witnessing problem in the associated theory. So the Σ^q_1 witnessing problem for G^*_i has the same complexity as the Σ^B_1 witnessing problem for V^i. This would be related to the work in [10], and may provide an alternative view of their results.

It would also be interesting to explore the provability of the reflection principles for formulas with low quantifier complexity. In particular, we could ask if V^i proves that G^*_{i+1} is consistent. People have thought of these problems ([7], Section 10.5), so the answers are not easy. However, this is still an important question to answer.

REFERENCES

[1] SAMUEL R. BUSS, *On Herbrand's theorem*, ***Lecture Notes in Computer Science***, vol. 960 (1995), pp. 195–209.

[2] S. A. Cook, *Feasibly constructive proofs and the propositional calculus*, ***Proceedings of the 7-th acm symposium on the theory of computation***, 1975, pp. 83–97.

[3] Stephen Cook, Theories for Complexity Classes and their Propositional Translations, Quaderni di Matematica, pp. 175–227, Quaderni di Matematica, 2003, pp. 175–227.

[4] Stephen Cook and Tsuyoshi Morioka, *Quantified propositional calculus and a second-order theory for NC^1*, ***Archive for Math. Logic***, vol. 44 (2005), no. 6, pp. 711–749.

[5] Stephen Cook and Phuong Nguyen, *Foundations of proof complexity: Bounded arithmetic and propositional translations*, Available from http://www.cs.toronto.edu/~sacook/csc2429h/book, 2006.

[6] Stephen Cook and Neil Thapen, *The strength of replacement in weak arithmetic*, ***ACM Trans. Comput. Logic***, vol. 7 (2006), no. 4, pp. 749–764.

[7] Jan Krajicek, ***Bounded arithmetic, propositional logic, and complexity theory***, Cambridge University Press, 1995.

[8] Jan Krajícek, Pavel Pudlák, and Gaisi Takeuti, *Bounded arithmetic and the polynomial hierarchy.*, ***Ann. Pure Appl. Logic***, vol. 52 (1991), no. 1-2, pp. 143–153.

[9] Jan Krajicek and Pavel Pulak, *Quantified propostitional calculi and fragments of bounded arithmetic*, ***Zeitschr. f. math. Logik und Grendlagen d. Math.***, vol. 36 (1990), pp. 29–46.

[10] Jan Krajicek, Alan Skelley, and Neil Thapen, *NP search problems in low fragments of bounded arithmetic*, ***The Journal of Symbolic Logic***, vol. 72(2) (2007), pp. 649–672.

[11] Alexis Maciel and Toniann Pitassi, *Conditional lower bound for a system of constant-depth proofs with modular connectives.*, ***Lics***, IEEE Computer Society, 2006, pp. 189–200.

[12] Tsuyoshi Morioka, *Logical approaches to the complexity of search problems: Proof complexity, quantified propositional calculus, and bounded arithmetic*, ***Ph.D. thesis***, University Of Toronto, 2005.

[13] Phuong Nguyen, *Separating dag-like and tree-like proof systems*, ***Lics '07: Proceedings of the 22nd annual ieee symposium on logic in computer science***, IEEE Computer Society, 2007, pp. 235–244.

[14] Steven James Perron, *Examining the fragments of G*, ***Lics '07: Proceedings of the 22nd annual ieee symposium on logic in computer science***, IEEE Computer Society, 2007, pp. 225–234.

UNIVERSITY OF TORONTO
DEPARTMENT OF COMPUTER SCIENCE,
M5S 3G4,
TORONTO, ONTARIO, CANADA
E-mail: sperron@cs.toronto.edu