Plausibly Hard Combinatorial Tautologies

Jeremy Avigad

Abstract. We present a simple propositional proof system which consists of a single axiom schema and a single rule, and use this system to construct a sequence of combinatorial tautologies that, when added to any Frege system, p-simulates extended-Frege systems.

1. Introduction

As was pointed out in [**6**], the conjecture that $NP \neq coNP$ can be construed as the assertion that there is no proof system (broadly interpreted) in which there are short (polynomial-length) proofs of every propositional tautology. Though showing $NP \neq coNP$ seems to be difficult, the above formulation suggests an obvious restriction, namely the assertion that specific proof systems are inefficient. One of the nicest results of this form to date is the fact that there are no short proofs of tautologies representing the pigeonhole principle in a fixed-depth Frege system (see, for example, [**1, 2**]). This approach to demonstrating a proof system's inefficiency seems natural: choose a suitable sequence of propositional formulas that express some true combinatorial assertion, and then show that these tautologies can't be proven efficiently by the system in question.

Unfortunately, in the case of Frege systems, there is a shortage of good candidates. Bonet, Buss, and Pitassi [**3**] consider a number of combinatorial principles with short extended-Frege proofs, and conclude that most of them seem to require at most quasipolynomial Frege proofs (see also [**7**] for further discussion). This isn't to say that there are *no* examples of tautologies whose Frege proofs are likely to require exponential length: Cook [**5**] has shown that propositional tautologies $Con_{EF}(n)$, which assert the partial consistency of extended-Frege systems, have polynomial extended-Frege proofs; whereas Buss [**4**] has shown that if a Frege system is augmented by these tautologies, it can polynomially simulate any extended-Frege system. In short, if there is any separation between Frege and extended-Frege proof systems, the assertions $Con_{EF}(n)$ witness this separation.

1991 *Mathematics Subject Classification.* Primary 03F20; Secondary 03B05.

Key words and phrases. proof complexity, propositional logic, Frege system, hard tautologies.

The obvious complaint is that such consistency assertions can hardly be called “combinatorial,” since they involve the coding of formulas, proofs, axioms, and rules of inference. However, the existence of such tautologies suggests that the problem is primarily one of esthetics: we have some plausibly hard tautologies, only they are not as natural as we would like them to be.

The point of this paper is to make the assertions $Con_{EF}(n)$ look a bit more combinatorial than they might seem at first glance. In the next section we present a surprisingly simple proof system, p-equivalent to any Frege system, which relies on a single logical connective, a single axiom, and a single rule. The consistency of this proof system then “translates” to somewhat combinatorial assertions regarding the hereditarily finite sets or directed acyclic graphs. It is well known that extended-Frege proof systems are essentially Frege systems in which formulas are represented by nodes of a circuit, and in Section 4 we show that the DAG formulation of the combinatorial assertion yields a sequence of tautologies that behave much like the propositions $Con_{EF}(n)$.

2. A Simple Proof System

We start by reviewing some definitions from [**6**] (see also [**8**]). A proof system F is said to be *implicationally complete* if whenever

$$\varphi_1, \varphi_2, \ldots, \varphi_n \models \psi$$

then ψ is derivable from $\varphi_1, \ldots, \varphi_n$ in F. A *Frege system* is an implicationally complete propositional proof system based on finitely many rules and axiom schemata. If F_1 and F_2 are two propositional proof systems then F_1 *p-simulates* F_2 if there is a polynomial time algorithm that translates any F_1 proof of a propositional formula ψ to an F_2 proof of the same formula. (This last definition can be modified to cover situations in which F_1 and F_2 are based on different sets of logical connectives.) In particular, if F_1 p-simulates F_2 then there is a polynomial bound to the increase in length of proof given by the translation.

The notion of a Frege system is extremely robust, in that any two Frege systems p-simulate each other, even if the underlying sets of connectives are different and connectives of variable arities are allowed. (This in stated in [**6**] and proven in [**10**]; for an alternative proof see [**4**].)

In this paper we will work with the single connective $nand(\varphi_1, \ldots, \varphi_k)$, where k is arbitrary. This connective asserts that not all of its arguments are true, or, equivalently, that at least one of them is false. It might be more accurate to think of *nand* as being a unary operator on sets, since, for example, we will treat $nand(\varphi_1, \varphi_2, \varphi_3)$, $nand(\varphi_3, \varphi_1, \varphi_2)$, and $nand(\varphi_1, \varphi_1, \varphi_2, \varphi_3)$ as the same formula. We will often use $\vec{\varphi}$ to denote a sequence of formulas $\varphi_1, \ldots, \varphi_k$.

It should be clear that *nand* is a complete connective, since one can define $\bot$ as $nand()$ and $\neg\varphi$ as $nand(\varphi)$, and then define the rest of the logical connectives in the usual way. Notice that propositions of the form

$$(Ax) \qquad nand(\vec{\psi}, \vec{\varphi}, nand(\vec{\varphi}))$$

are tautologically true, since if none of the φ_i are false, then $nand(\vec{\varphi})$ is. Similarly, the rule

$$(Cut) \qquad \frac{nand(\vec{\psi}, \vec{\varphi}) \quad nand(\vec{\psi}, nand(\vec{\varphi}))}{nand(\vec{\psi})}$$

is sound, since either all the φ_i are true, in which case the first premise guarantees the conclusion, or at least one of the φ_i is false, in which case $nand(\vec{\varphi})$ is true and the second premise guarantees the conclusion. Finally, the rule

$$(Weak) \qquad \frac{nand(\vec{\psi})}{nand(\vec{\psi}, \vec{\varphi})}$$

is sound, since the premise guarantees that at least one of the ψ_i is false. Notice that all the above rules are still valid even if the sets $\{\vec{\psi}\}$ and $\{\vec{\varphi}\}$ are not disjoint, or if either of these two sets is empty.

Let F_1 be the proof system consisting of (Ax), (Cut), and $(Weak)$, and let F_2 be the proof system consisting of just (Ax) and (Cut). We now have the following surprising

Theorem 2.1. *F_1 is a Frege-system, and F_2 p-simulates F_1.*

The proof of this theorem is subsumed by the following sequence of lemmas.

Lemma 2.2. *The rule*

$$(A) \qquad \frac{nand(\vec{\psi}, nand(\varphi_1)) \quad \ldots \quad nand(\vec{\psi}, nand(\varphi_l))}{nand(\vec{\psi}, nand(\varphi_1, \ldots, \varphi_l))}$$

is a derived rule of F_1.

Proof. Roughly speaking, rule (A) asserts that one can derive $\theta \vee (\varphi_1 \wedge \ldots \wedge \varphi_l)$ from $\theta \vee \varphi_1$ through $\theta \vee \varphi_l$. For simplicity let's consider the case $l = 2$, namely

$$\frac{nand(\vec{\psi}, nand(\varphi_1)) \quad nand(\vec{\psi}, nand(\varphi_2))}{nand(\vec{\psi}, nand(\varphi_1, \varphi_2))} .$$

Using an axiom, the first premise, weakening, and cut, conclude

$$\frac{nand(\vec{\psi}, \varphi_1, \varphi_2, nand(\varphi_1, \varphi_2)) \quad \dfrac{nand(\vec{\psi}, nand(\varphi_1))}{nand(\vec{\psi}, nand(\varphi_1), \varphi_2, nand(\varphi_1, \varphi_2))}}{nand(\vec{\psi}, \varphi_2, nand(\varphi_1, \varphi_2))} .$$

By weakening the second premise, conclude

$$\frac{nand(\vec{\psi}, nand(\varphi_2))}{nand(\vec{\psi}, nand(\varphi_2), nand(\varphi_1, \varphi_2))} .$$

The desired conclusion,

$$nand(\vec{\psi}, nand(\varphi_1, \varphi_2))$$

follows from a cut. The general case is similar and left to the reader. □

Lemma 2.3. *The rule*

$$(B) \qquad \frac{nand(\vec{\psi}, \vec{\varphi})}{nand(\vec{\psi}, nand(nand(\vec{\varphi})))}$$

is a derived rule of F_1.

Proof. Roughly speaking, rule (B) asserts that one can derive $\theta \vee (\varphi_1 \vee \ldots \vee \varphi_k)$ from $\theta \vee \varphi_1 \vee \ldots \vee \varphi_k$. By weakening the premise, conclude

$$nand(\vec{\psi}, \vec{\varphi}, nand(nand(\vec{\varphi}))).$$

The conclusion follows from a cut on the axiom

$$nand(\vec{\psi}, nand(\vec{\varphi}), nand(nand(\vec{\varphi}))).$$

□

Lemma 2.4. F_1 *is complete.*

Proof. Rules (A) and (B) are not only sound, but have the further property that if the conclusions are tautologies, then so are all the premises. The usual method of proving completeness by "working backwards" then applies: suppose α is a tautology of the form $nand(\vec{\varphi})$. If one of the φ_i is of the form $nand(\psi_1, \ldots, \psi_l)$ for $l > 1$ we can use the first rule to reduce the task of proving α to proving formulas in which the subformula φ_i is replaced by a unary *nand*. On the other hand, if one of the φ_i of the form $nand(\psi)$ and ψ is not a variable, then ψ must be of the form $nand(\vec{\theta})$ and we can use the second rule backwards to remove two layers of *nand*'s. Ultimately we are reduced to proving tautologies of the form $nand(\vec{\varphi})$, where each φ_i is either $nand()$, a variable, or the negation of a variable. But it is easy to see that such an assertion is a tautology iff either $nand()$ appears among the φ_i or some particular variable appears along with its negation, and in both cases the resulting tautology is an axiom. (Note that, in particular, $nand(nand())$ is an axiom.) □

Lemma 2.5. F_1 *is implicationally complete.*

Proof. Suppose $\varphi_1, \ldots, \varphi_k \models \psi$ where each φ_i is of the form $nand(\vec{\theta_i})$ and ψ is of the form $nand(\vec{\eta})$. Then

$$nand(\vec{\theta_1}) \wedge \ldots \wedge nand(\vec{\theta_k}) \rightarrow nand(\vec{\eta})$$

is a tautology. Some fiddling shows that this formula is equivalent to

$$nand(nand(\vec{\theta_1}), \ldots, nand(\vec{\theta_k}), \vec{\eta}),$$

and the fact that F_1 is complete implies that this is derivable in F_1. The conlusion $nand(\vec{\eta})$ then follows from the assumptions

$$nand(\vec{\theta}_1), \dots , nand(\vec{\theta}_k)$$

using k applications of weakening and cut. □

We can't quite say that F_2, which doesn't allow weakening, is implicationally complete; notice that in the previous lemma we had to weaken the assumptions $\varphi_1, \dots , \varphi_k$. However, we have the following

Lemma 2.6. *F_2 p-simulates F_1.*

Proof. We need to give a polynomial time algorithm that removes instances of weakening from any F_2 proof d. Since F_1 is a Frege system, we can assume that its derivations are tree-like (see [**7**]). Notice that

1. any two successive applications of weakening can be collapsed to a single one;
2. if φ follows from ψ_1 and ψ_2 by the cut rule, then any weakening of φ follows from the appropriate weakenings of ψ_1 and ψ_2; and
3. any weakening of an axiom is again an axiom.

As a result we can effectively eliminate instances of weakening in d, working from the bottom up. Notice that if the original proof d has length n, then it has at most n steps. The new proof will also have at most n steps, none of which is any longer than the length of d itself. A moment's reflection should convince the reader that, under the assumption that the original proof is tree-like, the algorithm can be made to run in polynomial time. □

In Section 4 we will need to assume that our proofs have a nice normal form.

Lemma 2.7. *Theorem 2.1 holds even if we assume that in the application of (Cut), $\{\vec{\varphi}\}$ is not a subset of $\{\vec{\psi}\}$, and $nand(\vec{\varphi})$ is not an element of $\{\vec{\psi}\}$.*[1]

Proof. Otherwise, the conclusion is the same as one of the hypotheses, in which case the proof tree can be pruned and this rule omitted. □

From now on we will assume that F_2 proofs are of this form. Notice that as a result, the arity of the conclusion of any cut rule is strictly less than the arity of either premise.

[1] The editor has pointed out that this lemma can be strengthened to require $\{\vec{\varphi}\}$ and $\{\vec{\psi}\}$ to be disjoint, albeit at the expense of a polynomial increase in the length of the proof.

3. A Theorem About Hereditarily Finite Sets

The hereditarily finite sets are defined inductively as follows: the empty set $\emptyset$ is a hereditarily finite set, and if $a_1, \ldots, a_k$ are hereditarily finite sets, then so is $\{a_1, \ldots, a_k\}$. There is a natural bijection between hereditarily finite sets and variable-free formulas built up with *nand*, given by the map f which takes sets $\{a_1, \ldots, a_k\}$ to formulas $nand(f(a_1), \ldots, f(a_k))$.

Andreas Blass has observed that every hereditarily finite set a corresponds to a two-player game, as follows. Player I starts by chosing any element $b \in a$, and player II must respond by chosing some element $c \in b$. Player I then plays some $d \in c$, and so on. The game continues until one player cannot move because the empty set has just been played, at which point this player has lost. The reader can verify that a closed formula φ is true if and only if Player I has a winning strategy in the corresponding game.

Formulas with variables can, of course, be identified with hereditarily finite sets over atoms $x_1, \ldots, x_k$. With this identification, the three rules from the previous section can be stated as follows:

1. Axiom: $a \cup b \cup \{b\}$
2. Cut: from $a \cup b$ and $a \cup \{b\}$ conclude a
3. Weakening: from a conclude $a \cup \{b\}$.

This observation allows us to cook up a somewhat combinatorial theorem about hereditarily finite sets.

Definition 3.1. *Say a hereditarily finite set c is* good *if it is of the form $a \cup b \cup \{b\}$; that is, there is some $b \in c$ such that $b \subset c$.*

So good sets correspond to axioms.

Theorem 3.2. *Let C be a hereditarily finite set, such that for every a in C, either*

1. *a is good, or*
2. *for some b that is neither an element nor a subset of a, $a \cup b$ and $a \cup \{b\}$ are both in C.*

Then the empty set is not in C.

Proof. By hypothesis, if a is in C then either a is good or a is "derived" from two strictly larger sets $a \cup b$ and $a \cup \{b\}$. Given any set a in C we can then work backwards and construct a "derivation" of a from good sets. But such a derivation of the empty set would then translate back to a proof of $\bot$ in F_2.

To phrase the proof slightly differently, let m be the maximum of the cardinalities of the sets in C, and show by induction on i that every set of cardinality $m - i$ in C corresponds to a true formula. But the empty set φ corresponds to $nand()$, which is false. □

The task of translating Theorem 3.2 into a sequence of propositional tautologies will be addressed in the next section. Roughly speaking, if we code hereditarily finite sets as strings, our tautologies will "express" the consistency of the Frege system F_1. On the other hand, to "express" the consistency of a corresponding *extended* Frege system, we will want to code hereditarily finite sets as nodes of a directed acyclic graph, as follows.

If G is a directed acyclic graph and a is a vertex of G, define the *neighborhood* of a to be

$$N(a) = \{b \mid \text{there is an edge from } a \text{ to } b\}.$$

We can think of each node a as coding a hereditarily finite set, consisting of the elements coded by the neighborhood of a. (If we wanted our representation to be canonical, we could demand that G be extensional, but this extra requirement is unnecessary for our purposes.)

Translating Theorem 3.2 to this new language yields the following

Theorem 3.3. *Let G be a directed acyclic graph, and suppose C is a subset of the vertices of C such that for every a in C, one of the following two conditions holds:*

1. *Either there is a vertex b in $N(a)$ such that $N(b) \subseteq N(a)$, or*
2. *there are vertices d and e in C, and a vertex b, such that*
   (a) *$N(d) = N(a) \cup \{b\}$,*
   (b) *$N(e) = N(a) \cup N(b)$,*
   (c) *$b \notin N(a)$*
   (d) *$N(e) \neq N(a)$.*

Then every element of C is nonterminal, that is, has at least one outgoing edge.

Proof. The last two conditions are the analogues of the requirements that b is neither an element nor a subset of a, and guarantees that the cardinalities of $N(e)$ and $N(d)$ are stricly larger than that of $N(a)$. As in the case of Theorem 3.2, a counterexample would unwind to a proof of a contradiction in F_1. □

Note that a graph G and subset C with satisfying the hypothesis of the theorem above could lead to an F_1 proof whose length is exponential in the size of the graph, since the DAG representation of formulas allows one to "reuse" subformulas efficiently. On the other hand, small graphs G *do* translate to small *extended*-Frege proofs. As it turns out, in the next section we will only need the converse of this fact, i.e. that small extended-Frege proofs translate to small graphs.

The following lemma provides a more direct proof of Theorem 3.3.

Lemma 3.4. *Let G be a directed acyclic graph. Then there is a set of vertices S such that for every a in G,*

1. *if a is in S, then $N(a) \cap \bar{S} \neq \emptyset$, and*
2. *if a is in $\bar{S}$, then $N(a) \subseteq S$.*

Proof. Intuitively, S consists of the nodes of G that correspond to true formulas. The construction of S proceeds in stages, just as in the assignment of truth values to the nodes of a circuit: we start by putting all the terminal nodes of G into $\bar{S}$, and once all the elements of $N(a)$ have been put into S or $\bar{S}$, we decide what to do with a based on clauses 1 and 2 of the lemma. □

The theorem then follows from the lemma, just as in the proof of Theorem 3.2, by noting that elements which satisfy clause 1 of the theorem are guaranteed to be in S, and if c and d are in S, so is any a satisyfing the condition set by clause 2. At the same time, no terminal node can be in S.

In the next section we show how to translate Theorem 3.3 to propositional tautologies that behave much like the tautologies $Con_{EF}(n)$. It would be nice if something resembling this theorem could be found somewhere in the graph theory literature. Though the odds are slim, the following reformulation might be more suggestive.

Theorem 3.5. *Let G be a directed graph, and suppose C is a set of nonterminal vertices of G satisfying the hypothesis of Theorem 3.3. Then G contains a cycle.*

4. Some Plausibly Hard Tautologies

Recall that an extended-Frege system (see [**6, 7**]) is obtained from a Frege system by allowing one to introduce constants A_φ at any point in a proof to abbreviate formulas φ, using the "extension axiom"

$$A_\varphi \leftrightarrow \varphi.$$

Note that φ may include other constants that have previously been introduced, and if "↔" is not among the basic connectives of the proof system, one can use any reasonable equivalent. The expectation is that these abbreviations "should" allow us to prove tautologies more efficiently, much the way that circuits "should" be able to represent boolean functions more efficiently than formulas. In fact, one can think of extended-Frege proofs as reasoning about formulas that are represented by nodes of a circuit, in a way we'll make explicit below.

To start with, let's fix an extended-Frege system EF_1, which consists of F_1 augmented by the extension axioms above. Define EF_2 to consist of F_2 augmented by axioms

$$nand(nand(A_\varphi), \varphi, \vec{\psi})$$

and

$$nand(nand(\varphi), A_\varphi, \vec{\psi}).$$

When added to F_1, these two axioms allow for short proofs of any weakening of an extension axiom of F_2, so just as in Theorem 2.1 we have that EF_2 p-simulates EF_1, and hence any extended-Frege system.

Given a proof in EF_1 (or EF_2), one can construct a directed acyclic graph that represents its formulas, in the following way.

Definition 4.1. *Let d be an EF_1- or EF_2-proof with variables x_1 to x_k, and let Γ denote the set of all subformulas of formulas in d. Say that a directed acyclic graph G with distinguished terminal nodes $\hat{x}_1, \dots, \hat{x}_k$* represents the subformulas of d *if there is a map h from Γ to G with the following properties:*

1. *$h(x_i) = \hat{x}_i$;*
2. *$h(nand(\varphi_1, \dots, \varphi_m))$ is a node a such that*
$$N(a) = \{h(\varphi_1), \dots, h(\varphi_n)\};$$
3. *$h(A_\varphi) = h(\varphi)$.*

If G represents the formulas of d via h, say that the subset

$$C = \{h(\varphi) \mid \varphi \text{ is a line of } d\}$$

represents the proof d *in G.*

It should be clear that for any proof d we can find a G representing its subformulas, such that the size of G is polynomial in the length of d. Notice that h maps extension axioms of EF_2 to "good" nodes, that is, nodes satisfying Clause 1 of Theorem 3.3.

Now it is not difficult to cook up, for each n, a tautology that expresses Theorem 3.3 for graphs of size n, such that the length of these tautologies is bounded by a polynomial in n. We can use variables p_{ij} for $i < j$ to represent the assertion that there is an edge from i to j (the condition $i < j$ guarantees that the graph is acyclic), and variables q_i to express the assertion that i is in C. The hypothesis of the theorem is then a conjunction

$$\bigwedge_i (q_i \to \varphi_1(i) \lor \varphi_2(i))$$

where $\varphi_1(i)$ is the assertion

$$\bigvee_j \left(p_{ij} \land \bigwedge_k (p_{jk} \to p_{ik}) \right)$$

and $\varphi_2(i)$ is the assertion

$$\bigvee_{j,k,l,m} \left(q_k \land q_l \land p_{kj} \land \bigwedge_{n \neq j} (p_{kn} \leftrightarrow p_{in}) \land \bigwedge_n (p_{ln} \leftrightarrow (p_{in} \lor p_{jn})) \land \neg p_{il} \land \neg p_{im} \land p_{jm} \right).$$

Here all the variables in the large conjunctions and disjunctions range from 1 to n, and by p_{ij} for $i \geq j$ we really mean $\bot$. The conclusion of the theorem, that is, the assertion that there is no terminal edge in C, translates to

$$\bigwedge_i (q_i \to \bigvee_j p_{ij}).$$

Call the resulting tautology $T(n)$.

We now come to the main theorems in this section.

Theorem 4.2. *The tautologies $T(n)$ have polynomial-size extended-Frege proofs.*

Proof. This can be proven in much the same way that one proves the equivalent result for $Con_{EF}(n)$, as in [**5**] or [**7**]; the set S of Lemma 3.4 can be defined in PV or V_1^1. □

Theorem 4.3. *Let F be any Frege system, and let $\hat{F}$ be the proof system obtained by additonally allowing any substitution instance of a tautlogy $T(n)$. Then $\hat{F}$ p-simulates any extended-Frege system.*

Proof. Our proof is modeled on the equivalent result for $Con_{EF}(n)$, which appears in [**4**, Main Theorem 3]. Both arguments rely heavily on the fact that Frege systems have an adequate formalization of the notion for truth for propositional formulas, which is the main innovation in the paper just cited.[2]

Given an extended-Frege proof d of $\varphi(x_1, \ldots, x_n)$, we need to show how to construct a proof of $\varphi(x_1, \ldots, x_n)$ in $\hat{F}$, whose length is polynomially bounded in the length of d.

Since EF_2 p-simulates any extended-Frege system, we can efficiently construct, from d, a "DAG proof" of $nand(nand(\varphi))$; that is, a directed acyclic graph G and a set C representing an EF_2 proof of that formula.

Given any assignment E of truth-values (that is, constants $nand()$ and $nand(nand())$ representing false and true, respectively) to the variables $x_1, \ldots, x_n$, let φ^E denote the closed formula $\varphi(E(\vec{x}))$. We now argue in F. Letting E be arbitrary, we can replace each node $\hat{x}_i$ of G by a graph representing $E(x_i)$, and obtain a DAG proof of $nand(nand(\varphi^E))$. On the other hand, if φ evaluated at E is false, we can construct (as in [**4**]) a DAG proof of $nand(\varphi^E)$; more precisely, there is a sequence of formulas defining a DAG proof of this formula, whose length is polynomial in the length of d, such that F can verify that this DAG proof satisfies the hypothesis of Theorem 3.3.[3] Gluing these two together with a cut, we obtain DAG proof of $nand()$, that is, a counterexample to Theorem 3.3.

In short, we can show with a polynomial-size proof in F that if φ^E is false, then a substitution instance of one of the tautologies $T(n)$ fails. As a result, $\hat{F}$ proves that φ is true under an arbitrary assignment of truth values to its variables. By the adequacy of the truth predicate, $\hat{F}$ can then conclude φ. □

[2] For the argument below, one can either adapt the truth predicates of [**4**] to handle *nand*'s of arbitrary arity, or just work with their infix equivalents in F.

[3] The argument is actually a bit more delicate than the one outlined in [**4**]. Since EF_2 doesn't allow weakening, we have to "anticipate" the weakening rules that we would have used in the proof described there. Working on $nand(\varphi)$ from the "top down," we can assign to each suformula ψ appropriate sequences $\vec{\theta}_{\psi,1}$ and $\vec{\theta}_{\psi,2}$; then, from the "bottom up," prove the following in F: "if ψ^E is true, there is a DAG proof of $nand(nand(\psi^E), \vec{\theta}^E_{\psi,1})$, and if ψ^E is false, there is a DAG proof of $nand(\psi^E, \vec{\theta}^E_{\psi,2})$."

COROLLARY 4.4. *If there is a superpolynomial separation between Frege systems and extended-Frege systems, then there are no polynomial size proofs of the tautologies $T(n)$ in any Frege system.*

PROOF. If a Frege system F had polynomial-size proofs of the tautologies $T(n)$, then substituting formulas for variables would yield polynomial-size proofs of arbitrary substitution instances of these tautologies. By the previous theorem, F would then p-simulate any extended-Frege system. □

5. Final Comments

The system F_2 is not the first axiomatization of propositional logic based on a single axiom and rule. In 1913 Sheffer showed that the binary nand "$\varphi \mid \psi$" is a complete connective, and soon after, Jean Nicod [**9**] showed that the axiom schema

$$\{[p \mid (q \mid r)] \mid [t \mid (t \mid t)]\} \mid \{[s \mid q] \mid [(p \mid s) \mid (p \mid s)]\}$$

combined with the rule

$$\frac{p \mid (r \mid q) \quad p}{q}$$

provide a complete axiomatization of propositional logic. (Here "complete" means that one can derive the axioms and rules of Russell and Whitehead's *Principia Mathematica*; the modern notion of completeness for propositional logic first appeared independently in the work of Bernays and Post, a few years later.) In the second edition of the *Principia*, which appeared in 1925, the authors cite this work, and, oddly enough, call Sheffer's reduction "the most definite improvement resulting from work in mathematical logic during the past fourteen years."

The fact that the proof system F_2 presented here p-simulates any Frege system tells us that without loss of generality we can think of any Frege proof as a tree, with a tautology at the bottom, cuts at the branches, and axioms ("the good sets") at the nodes. If would be nice if this simple formulation could be put to good use in proving lower bounds.

References

[1] Ajtai, Miklos, "The complexity of the pigeonhole principle," *Proceedings of the IEEE 29th Annual Symposium on Foundations of Computer Science* (1988), pp. 346-355.

[2] Beame, Paul, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, Pavel Pudlák, and Alan Woods, "Exponential lower bounds for the pigeonhole principle," *Proceedings of the 24th Annual ACM Symposium on Theory of Computing* (1992), pp. 200-221.

[3] Bonet, Maria, Samuel Buss, and Toniann Pitassi, "Are there hard examples for Frege systems?" in Clote and Remmel eds., *Feasible Arithmetic II*, Birkhauser, 1995.

[4] Buss, Samuel, "Propositional consistency proofs," *Annals of Pure and Applied Logic*, vol. 52 (1991), pp. 3-29.

[5] Cook, Stephen, "Feasibly constructive proofs and the propositional calculus," *Proceedings of the 7th Annual ACM Symposium on Theory of Computing* (1975), pp. 83-97.

[6] Cook, Stephen and Robert Reckhow, "The relative efficiency of propositional proof systems," *Journal of Symbolic Logic*, vol. 44 (1979), pp. 36-50.

[7] Krajíček, Jan, "On Frege and extended-Frege proof systems," in Clote and Remmel eds., *Feasible Arithmetic II*, Birkhauser, 1995.

[8] Krajíček, Jan, *Bounded Arithmetic, Propositional Logic, and Complexity Theory,* Cambridge University, 1995.

[9] Nicod, Jean, "A reduction in the number of the primitive propositions of logic," Proceedings of the Cambridge Philosophical Society, vol. 19 (1917), pp. 32-41.

[10] Reckhow, Robert, *On the Lengths of Proofs in the Propositional Calculus*, Ph. D. thesis, University of Toronto, 1976.

Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA 15213

E-mail address: `avigad+@cmu.edu`