Alcom-FT Technical Report Series
ALCOMFT-TR-02-132

The Complexity of Resource-Bounded Propositional Proofs

Doctoral Thesis presented to the
Departament de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya
by
Albert Atserias
atserias@lsi.upc.es

December 18, 2001

Co-directed by
José Luis Balcázar
María Luisa Bonet

Abstract

Propositional Proof Complexity is an active area of research whose main focus is the study of the length of proofs in propositional logic. There are several motivations for such a study, the main of which is probably its connection to the P vs NP problem in Computational Complexity.

The experience in the field has revealed that the most interesting parameter of a proof system is the set of allowed formulas. The exact set of rules and axioms of the system is often irrelevant as long as they remain sound and reasonable. This parameterization by the set of allowed formulas establishes a link with the field of Boolean complexity where bounds are imposed on different computational resources as a classification tool. This idea is adopted in the proof complexity setting too. In this thesis we study the complexity of several 'resource-bounded' proof systems.

The first set of results of the thesis is about a resource-bounded proof system that we call the Monotone Sequent Calculus (MLK). This is the standard propositional Gentzen Calculus (LK) when negation is not allowed in the formulas. The main result is that the use of negation does not yield exponential savings in the length of proofs. More precisely, we show that MLK quasipolynomially simulates LK on monotone sequents. We also show that, as refutation systems, MLK is polynomially bounded if and only if LK is. These results are in sharp contrast with the situation in Boolean complexity where the use of negations provably yields an exponential gap in computational power.

The second set of results is about a proof system extending Resolution by allowing disjunctions of conjunctions of two literals and not only disjunctions of literals. We prove that the Weak Pigeonhole Principle with twice as many pigeons than holes, and random 3-CNF formulas, require exponential-size proofs in this system. The interest of our results is that we break the barrier of one alternation of disjunctions and conjunctions in the context of the weak pigeonhole principle and random 3-CNF formulas for the first time. The techniques that we employ combine old ideas with new arguments using powerful tools from probability theory.

The third and final set of results is about the exact complexity of the Weak Pigeonhole Principle with twice as many pigeons than holes. We show that all previously known quasipolynomial-size proofs of bounded-depth are not optimal in terms of size. As a consequence of our result, we obtain Bounded Arithmetic proofs that there are infinitely many primes from 'large number assumptions' that are provably weaker than previously needed.

Acknowledgments

I want to thank all those who made this thesis possible. First, I want to thank my advisors José Luis Balcázar and María Luisa Bonet for accepting co-directing this work. Their accurate technical advise and permanent encouragement made possible the progress reported here.

I am also in debt with Jorge Castro, Victor Dalmau, Juan Luis Esteban, Nicola Galesi and David Guijarro for the good times we had while discussing the (sometimes crazy) ideas that led (but sometimes not) to the results. Special thanks go to Juan Luis Esteban, Nicola Galesi and Pavel Pudlák who allowed me to include our joint work as contributions of the thesis. Thanks also to the CUR of the Generalitat de Catalunya who supported me through grant 1999FI 00532. Some of the trips were also supported by ALCOM-FT [1] and FRESCO [2] projects.

Finally, I want to mention my family and friends for their help, and Sandra for *every thing*. I dedicate this thesis to her.

[1]Partially supported by the Future and Emerging Technologies programme of the EU under contract number IST-1999-14186 (ALCOM-FT).

[2]Partially supported by MEC PB98-0937-C04 (FRESCO).

Contents

Chapter 1

Introduction

1.1 Background

Since Boole turned logic into a mathematical discipline, its influence in modern mathematics has been tremendous. Hilbert identified Mathematical Logic as the tool with which to establish the foundations of mathematics through his well-known Consistency Program. Unfortunately, Gödel's Incompleteness Theorems showed that the program could not work. Nonetheless, the failure of Hilbert's program was reinterpreted by logicians in their own benefit: although mathematics cannot be mechanized, logic is the tool to approximate that goal as much as possible. A first step towards this is understanding the incompleteness phenomenon in its full strength. In a sense, Computability Theory is the result of such a quest.

With the possibility of building actual computers, Computability Theory was refined into a number of different subdisciplines. Among them is Complexity Theory, whose development shares a number of similarities with the history of the consistency program. Von Neumann, among others, turned Boolean logic into the calculus of actual computers. Later on, with the empirical observation that many algorithmic problems seemed intractable in real life, Cook identified logic, again, as the tool with which to establish the foundations of computer intractability. Cook's conjecture that $\mathbf{P} \neq \mathbf{NP}$ is often interpreted as an incompleteness phenomenon down to real life computability. Needless to say, any form of logic, be it (Feasible) Proof Theory, (Finite) Model Theory, or (Efficient) Computability Theory, plays a fundamental role in today's research towards $\mathbf{P} \neq \mathbf{NP}$. We focus here in Feasible Proof Theory, also called Proof Complexity.

1.2 Propositional Proof Complexity

Proof systems for propositional logic should be *sound*, which means that every provable formula should be a tautology, and *complete*, which means that every tautology should be provable. Moreover, it is desirable that proofs be *easy to check* and, if possible, *short*. While most usual proof systems appearing in textbooks satisfy the first three requirements, it is not clear at all whether they satisfy the condition that every tautology has a short proof. Informally, the aim of Propositional Proof Complexity is to study this phenomenon mathematically.

Following Cook and Reckhow [CR79], a propositional proof system $\mathcal{P}$ is a polynomial-time computable function from the set of words over a finite alphabet Σ^* onto TAUT, the set of propositional tautologies reasonably encoded in a finite alphabet. If $\mathcal{P}(x) = F$ for $x \in \Sigma^*$ and $F \in \mathrm{TAUT}$, we say that x is a proof of F in $\mathcal{P}$. From its very definition, we see that a propositional proof system is sound, complete, with easy to check proofs: soundness follows from the fact that its range is the set of tautologies, completeness follows from the fact that the function is onto the set of tautologies, and proofs are easy to check since the function is polynomial-time computable. Moreover, all usual proof systems for propositional logic such as Resolution, Hilbert Style Axiomatic Systems, or Gentzen Calculus fit into this definition. As mentioned before, the aim of Propositional Proof Complexity is to study *the length of proofs*. There are at least three sources of motivation for such an enterprise.

Cook's Program: towards $\mathbf{NP} \neq \mathbf{coNP}$. The goal of Computational Complexity Theory is to classify algorithmic problems according to the amount of resources that are required to compute them on formal models of computation. In a seminal paper, Cook [Coo71] identified $\mathbf{P}$ and $\mathbf{NP}$ as fundamental complexity classes delineating the boundary between computer tractability and intractability. The main result of Cook's work implied that TAUT does not belong to $\mathbf{P}$ unless $\mathbf{P} = \mathbf{NP}$. Whether $\mathbf{P} = \mathbf{NP}$ has become one of the most challenging problems in mathematics [Coo00].

In view of Cook's results above, it is not surprising that the complexities of proofs and computations be tightly related. In a later article, Cook and Reckhow [CR79] proved that $\mathbf{NP} \neq \mathbf{coNP}$ if and only if no propositional proof system is polynomially bounded, that is, if and only if for every propositional proof system $\mathcal{P}$ there exists a family of tautologies whose shortest proofs in $\mathcal{P}$ are not bounded by a polynomial in the size of the tautologies. Thus, if we identify *short* with *polynomially bounded*, we see that the question of whether propositional proof systems always have *short* proofs is a fundamental one.

Clearly, proving that increasingly powerful proof systems are not polynomially bounded is partial progress towards proving $\mathbf{NP} \neq \mathbf{coNP}$. This approach is known as *Cook's Program* and has been adopted by a number of researchers in recent years.

Feasible Mathematics. While the fundamental question is: Is it provable?, the practical questions are: Does it have a short proof? Do we need all that machinery to prove it? Although we classify these questions as practical, truly fundamental questions are hidden in them. Indeed, the more elementary the machinery we use is, the more constructive the proof usually becomes. As a slogan, constructive proofs tell us more than simply the truth. For example, a truly feasible proof that every number has a prime factor would give an (as yet unknown) efficient algorithm to factor large numbers, and so cryptography would not be possible as conceived today. Of course, this very remark can be turned upside down: knowing that certain statements do not have feasible proofs is also valuable information.

There are tight connections between provability in weak theories of arithmetic and the lengths of proofs in a propositional setting. Thus, proving that certain propositional tautologies do not have short proofs implies that certain arithmetical statements do not have feasible proofs. Although the converse is not always true, in cases of interest it usually is. There is a particularly interesting case that has been considered recently by Razborov [Raz95b, Raz95a, Raz96, RR97]. This concerns the metamathematics of fundamental questions such as $\mathbf{P} \neq \mathbf{NP}$. Razborov proves a certain degree of unfeasibility of its proof (if it can be proved) assuming a widely believed cryptographic conjecture.

Algorithmic Aspects of Proof Complexity. Automated Theorem Provers implement algorithms that find proofs in particular proof systems. Their actual running complexity is often unknown since they implement sophisticated heuristics to guide the search. However, if one is able to show that a tautology of interest requires large proofs in the underlying proof system, that is a proof of the inefficiency of the theorem prover since the output itself is already large. Hence, the results of proof complexity serve as a tool of analysis, as well as a source of test cases.

On the other hand, one may ask a different kind of question regarding proof systems. For example, one may be interested in efficiently finding short proofs only when such proofs exist. A proof system that admits an algorithm to perform such a task is called automatizable. It turns out that in an attempt to prove a lower bound for a certain proof system, one may find a battery of interesting properties that may show useful to develop theorem provers. The most prominent example is a weak form of automatizability of Resolution that was discovered recently by Ben-Sasson and Wigderson [BSW01].

1.3 Resource-Bounded Proof Systems

As mentioned in Section 1.2, proving strong lower bounds for arbitrary proof systems is a task at least as difficult as proving that $\mathbf{P} \neq \mathbf{NP}$. Therefore, similar to the approach taken in the field of Boolean Circuit Complexity where severe restrictions are imposed on circuits in order to prove lower bounds, it is natural to focus on restricted classes of proof systems. Although Cook's program explicitly proposes studying well-known proof systems, the experience in the field has revealed that the most interesting parameter of a proof system is the class of allowed formulas. It is often the case that the particular set of axioms and rules is not important as long as they remain sound, complete and reasonable. In this thesis, we will be mainly interested in studying the effect of restricting the class of allowed formulas. A different sort of restriction that one can impose on a proof system is requiring some particular structural form in the derivation, such as being a tree. This approach has been widely studied and will not be adopted here (see [Gal00]).

The less stringent restriction on the type of formulas is allowing arbitrary Boolean circuits. Typical proof systems using such a class of formulas are the Extended Frege systems. An apparently stronger restriction is requiring that the allowed formulas of the system be true Boolean formulas, that is, Boolean circuits of fan-out one (each gate feeds no more than one gate). Again, typical proof systems of this type are Frege Systems and the propositional Gentzen Calculus. Probably, the most stringent restriction of interest is requiring the allowed formulas to be clauses, that is, disjunctions of literals. The typical example of a proof system with such a requirement is Resolution. We note that in order to deal with arbitrarily complex tautologies, even if stringent restrictions are imposed on the allowed formulas, we often consider the refutational version of the proof system, in which the negation of the tautology is represented as a set of contradictory clauses, and the goal is to derive a contradiction from them.

Between Extended Frege systems and Resolution, there is a whole spectrum of *resource-bounded* proof systems. One may impose formulas with a limited number of alternations of conjunctions and disjunctions (the so-called depth of the formula), a limited use of some Boolean connective (such as negation), a limited size measured by the number of distinct literals that occur, a limited degree when expressed as a polynomial over some ground field, or any combination of these and other possibilities thereof.

There are several restrictions that are motivated from the field of Boolean Circuit Complexity. Restricting the depth of Boolean formulas is particularly interesting. On the one hand, this restriction generalizes Resolution whose allowed formulas have depth 0. On the other hand, the current understanding of bounded-depth circuits led to the discovery of quite spectacular lower bounds in Proof Complexity. Moreover, it has turned out that proof systems with formulas of bounded-depth are, in a precise technical sense, tightly related to weak theories of arithmetic such as $I\Delta_0$. We will work with these proof systems and weak theories of arithmetic in this thesis.

A second example of particular interest is restricting the use of negations in formulas. In the realm of Boolean Circuit Complexity, restrictions of these sort have lead to the most successful lower bounds. In Proof Complexity, there is a nice approach towards introducing such monotonicity restrictions. Pudlák [Pud99] proposed studying the propositional Gentzen Calculus without negation. The proof system is also called Geometric Logic. It turns out that several tautologies of interest can be written as sequents without negation. The typical example is the Pigeonhole Principle which can be expressed as follows. If variable $p_{i,j}$ means that pigeon i sits in hole j, the sequent says that if $n+1$ pigeon sit in n holes, some hole must contain two different pigeons. In symbols:

$$\bigwedge_{i=1}^{n+1} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i \neq j}}^{n+1} p_{i,k} \wedge p_{j,k}.$$

This proof system will also be thoroughly studied in this thesis.

Finally, we want to warn the reader that the so-called *algebraic proof systems* will not be considered in this thesis. Those systems have quite a different flavor. In particular, the allowed set of axioms and rules often make a difference there. Needless to say, the type of problems and techniques one encounters with these systems are of primary interest too (see [Pit97] for a survey).

1.4 Some Interesting Tautologies

1.4.1 The Weak Pigeonhole Principle

Informally, the Pigeonhole Principle (PHP) states that if $n+1$ pigeons sit in n holes, then some hole must contain at least two pigeons. In general, and more formally, the PHP_n^m with $m > n$ states that there is no one-to-one mapping from a set of m elements into a set of n elements. When m is significantly larger than n (say $m \geq 2n$), the principle is called the Weak Pigeonhole Principle (WPHP). Although one may tend to think that it is a trivial statement, the PHP states a fundamental property of cardinality of sets and appears recurrently in mathematical arguments, notably in combinatorics, probability theory, and number theory. To begin with, the induction principle is a special case of the pigeonhole principle. We present the argument in [Kra95] since we find it instructive. Suppose that $P(x)$ is a property of the natural numbers that violates the induction principle, that is, suppose that there exists a natural number n such that $\neg P(n)$ but $P(0) \wedge (\forall x)(P(x) \rightarrow P(x+1))$ holds. Then, it is easy to see [1] that the function

$$f(x) = \begin{cases} x & \text{if } P(x) \\ x-1 & \text{if } \neg P(x) \end{cases}$$

maps $\{0,\ldots,n\}$ one-to-one into $\{0,\ldots,n-1\}$. This violates the pigeonhole principle and we are done.

WPHP in Mathematics and the Theory of Computation. The WPHP is a combinatorial statement expressing a basic Ramsey-type property of mappings: any mapping from a large set into a small set contains unavoidable patterns (a collision). Therefore, we expect arguments in Finite

[1]Indeed, suppose that $f(a) = f(b)$ but $a \neq b$. Necessarily, either $a = b+1$ or $b = a+1$ since $f(a) = f(b)$. Without loss of generality, $a = b+1$. Now $f(b) = f(a) \geq a-1 = b$ and so $f(b) = b$ since $f(b) \leq b$. Hence, $P(b)$ holds and so does $P(b+1)$ since $(\forall x)(P(x) \rightarrow P(x+1))$. Therefore $f(a) = f(b+1) = b+1 = a$. We conclude that $f(a) \neq f(b)$; contradiction. Note that this argument is really elementary (it could be formalized in $Iopen$, say).

Combinatorics to make use of the WPHP more or less explicitly. As a matter of fact, counting arguments usually do, and these are everyday's tools in combinatorics.

The Probabilistic Method was introduced by Erdös as a method to prove non-emptiness of sets. The method consists in assigning a probability measure on sets of interest, and prove that our particular set has non-null measure. In finite combinatorics, the probabilistic method reduces to a counting argument, which in turn reduces typically to the WPHP. As a second example, the lower bounds on the size of circuits are typically counting arguments. The clearest example of this fact is Shannon's lower bound for almost all Boolean functions [Sha49]: the argument formalizes the observation that there are many more Boolean functions on n variables than small circuits with n inputs. Similarly, the new proof of Hastad's Switching Lemma due to Razborov [Raz95a] is a counting argument that can be formalized using the WPHP. For a more indirect and sophisticated use of the WPHP, we mention Paris, Wilkie and Woods [PWW88] feasible proof of the infinitude of primes: they find a proof in the bounded arithmetic theory $I\Delta_0$ that there are infinitely many primes assuming the PHP^{2n}_n as an axiom for every $n > 0$.

There are also some connections between the complexity of proving the WPHP and the hardness of computational problems. Krajícek, Pudlák and Takeuti [KPT91] proved that if $\mathbf{NP} \not\subseteq \mathbf{P/poly}$, then the bounded arithmetic theory PV does not prove the PHP^{2n}_n. More recently, Krajícek and Pudlák [KP95] proved that if the cryptosystem RSA is secure against polynomial attacks, then the bounded arithmetic theory S^1_2 does not prove PHP^{2n}_n. We note, by the way, that a result of Paris, Wilkie and Woods [PWW88] as refined by Krajícek [Kra95] implies that PHP^{2n}_n is provable in S^3_2.

Finally, Razborov has connected the complexity of proving WPHP with the metamathematics of fundamental problems in Complexity Theory such as $\mathbf{P} \neq \mathbf{NP}$. Indeed, Razborov argued that the problem $\mathbf{P} \neq \mathbf{NP}$ can encode a certain WPHP [Raz98]. Hence, proving that such a WPHP requires unfeasible proofs would be telling us that any proof of $\mathbf{P} \neq \mathbf{NP}$ must be unfeasible.

WPHP in Proof Complexity. In the setting of propositional proofs, the complexity of the Weak Pigeonhole Principle has been and remains to be a challenge for the community. In a sense, it is one of those classical problems whose partial solutions have stimulated the advance of the field with the introduction of new methods and techniques. In the context of Resolution, its exact complexity has remained open until very recently, and that contributed to the introduction of new methods. Here is a short account of this history. Haken [Hak85] proved an exponential lower bound for the size of Resolution proofs of the Pigeonhole Principle PHP^{n+1}_n. Later, Buss and Turán [BT88] extended his technique to obtain exponential lower bounds for PHP^m_n when $m = o(n^2/\log(n))$. The complexity of $\mathrm{PHP}^{n^2}_n$ in Resolution remained widely open: the method failed badly. This indicated that, in spite of having already proved exponential lower bounds for Resolution, the strength of the proof system was not completely understood yet. The introduction of the Width

Method for Resolution by Ben-Sasson and Wigderson [BSW01], following the work of Beame and Pitassi [BP96], shed some more light to the understanding of Resolution. Unfortunately, the width method still failed for a number of tautologies including the $\mathrm{PHP}_n^{n^2}$. The work of Bonet and Galesi [BG99], showing the intrinsic limitations of the width method, ratifyied this point.

Researchers started at least three new lines of attack. The first line of attack consisted in restricting the form of Resolution proofs in order to gain understanding on the problem. Tree Resolution was considered by Buss and Pitassi [BP97] who proved exponential lower bounds for PHP_n^m. This lower bound can also be obtained through the width method, though. A notion of rectangular Resolution proof was introduced by Razborov, Wigderson and Yao [RWY97], and exponential lower bounds were also proved. Regular Resolution was considered by Pitassi and Raz [PR01] who proved exponential lower bounds too. Each of these papers introduced some new ingredient to the final solution of the problem which is of independent interest. The second line of attack consisted in looking at tautologies sharing some of the properties of the Weak Pigeonhole Principle. The typical example is the Mutilated Chess Board Tautology stating that a squared grid with two opposite corners cut off cannot be covered by non-overlapping pieces each covering two adjacent squares of the grid. The difficulty in proving lower bounds for this tautology was, again, the failure of the width method. The problem was solved by Aleknovich [Ale00] and Dantchev and Riis [DR00] independently. This result, while being progress towards the problem on the Weak Pigeonhole Principle, is interesting in its own right. The third line of attack was suggested by Krajicek. He realized that exponential lower bounds for $\mathrm{PHP}_n^{n^{3/2}}$ in a system stronger than Resolution would imply exponential lower bounds for $\mathrm{PHP}_n^{n^2}$ in Resolution. This system, called Res(2) by Krajicek, consists in augmenting the set of allowed formulas to disjunctions of conjunctions of two literals. He suggested studying the Res(2)-systems by themselves, and that motivated the work and the new techniques of Chapter 4 of this thesis. As expected, the interest of these result will go beyond the problem of the Weak Pigeonhole Principle. Finally, Raz [Raz01a] completely solved the problem by proving an exponential lower bound for Resolution proofs of PHP_n^m for any $m > n$. Later Razborov [Raz01b] improved Raz's results in various aspects.

In view of this, we can say that it has been quite a successful approach for the field to have the Weak Pigeonhole Principle as a challenge. Although the problem has been finally solved for Resolution, there is still the challenge of knowing its exact complexity in systems with formulas of bounded-depth. Part of the work of this thesis is motivated by this problem, and we will return to it later in this introduction.

1.4.2 Random CNF Formulas

We already pointed out in Section 1.3 that tautologies will often be represented by their negation, as contradictory sets of clauses, or as we say, unsatisfiable CNF formulas. It turns out that researchers of various fields have been interested in defining a probability distribution on CNF formulas analogous to the random graph model of Erdös and Renyi [ER60]. The goal is, as usual, to gain understanding on the structure of these complex combinatorial objects. However, some researchers like to make emphasis on the interest of studying the complexity of the satisfiability problem on the average case.

The most commonly used probability distribution on CNF formulas is defined as follows. For a fixed number of propositional variables n, a fixed number of clauses m, and a fixed number of literals per clause k, a random k-CNF formula is produced by choosing exactly m clauses, each drawn independently from the uniform distribution on clauses of exactly k literals on distinct variables. This probability distribution is denoted $\mathcal{F}_n^{m,k}$. Several empirical and theoretical results are known about these distributions. A crucial parameter is the ratio of clauses to variables $\Delta = m/n$. The most basic and interesting fact is the observation that when Δ is below a certain threshold θ_k, the CNF formula is almost surely satisfiable, and when Δ is above θ_k, the formula is almost surely unsatisfiable. Experimental results suggest that $\theta_3 \approx 4.2$, while theoretical results can prove that $3.03 \leq \theta_3 \leq 4.7$ if θ_3 exists at all [CS88, CF90, CR92, FS90, KKKS98].

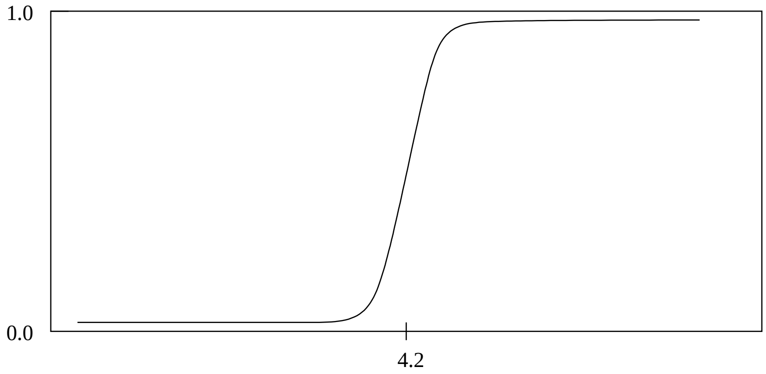

Figure 1.1: The probability that a random 3-CNF is unsatisfiable as a function of the ratio $\Delta = m/n$.

Interestingly enough, there is an observed degree of difficulty in deciding satisfiability when Δ is near the threshold. This empirical result was later supported by a result of Chvátal and Szeméredi [CS88] saying that for Δ near the threshold, every Resolution refutation of a random 3-CNF formula is almost surely of size $2^{\Omega(n)}$. These results were later extended by Beame and Pitassi [BP96], and Beame, Karp, Pitassi and Saks [BKPS98]. At the time of writing, it is not known if random 3-CNF formulas are almost surely hard to refute in arbitrary strong proof systems. One of the results of this thesis is proving lower bounds for random 3-CNF formulas in a proof system strictly stronger than Resolution.

1.5 Overview of the Results of the Thesis

The goal of this thesis is to study the size of propositional proofs when the allowed formulas have some specified bound on a resource. The two primary resources of interest are (1) negations and (2) alternations of conjunctions and disjunctions. We elaborate a little bit more on the motivations while we discuss the results.

The Monotone Calculus Boolean circuits that do not use negations are called Monotone. This is motivated by the fact that such circuits compute Boolean functions that are monotone with respect to the componentwise partial order. More formally, a Boolean function $f : \{0,1\}^n \to \{0,1\}$ is called monotone if for every $x = x_1 \ldots x_n \in \{0,1\}^n$ and $y = y_1 \ldots y_n \in \{0,1\}^n$, if $x_i \leq y_i$ for every $i \in \{1, \ldots, n\}$, then $f(x) \leq f(y)$. It is easy to see that every Boolean circuit over the basis $\{\wedge, \vee, 0, 1\}$ computes a monotone Boolean function, and that every monotone Boolean function is computed by a Boolean circuit over $\{\wedge, \vee, 0, 1\}$. For this reason, such circuits are called Monotone Circuits.

Since many interesting Boolean functions are monotone, it is natural to ask for their monotone complexity, that is, the size of the minimal monotone circuits computing them. Several lower bounds are known for this model. The most important of all is Razborov's [Raz85a] lower bound for the size of monotone circuits computing the Clique function, as improved by Alon and Boppana [AB87]. The same techniques led to the proof that the use of negation can make a difference for computing monotone functions: Razborov [Raz85b] proved that the Perfect Matching function requires quasipolynomial-size monotone circuits while polynomial-size circuits (with negation) exist for this function. The gap between monotone and non-monotone circuits was later improved to an exponential by Tardos [Tar87] for a different function. Monotone formulas, that is, Boolean formulas over the basis $\{\wedge, \vee, 0, 1\}$, have also been considered and several lower bounds are known [KW90, RW92].

Since the restriction to the monotone basis $\{\wedge, \vee, 0, 1\}$ has been so fruitful in proving lower bounds, Pudlák [Pud99] suggested a similar approach in Proof Complexity. He proposed studying the propositional Monotone Gentzen Calculus MLK which is obtained from the standard propositional Gentzen Calculus LK when negations are not allowed (see Section 1.3). In view of the results mentioned above for monotone circuit complexity, one would expect an exponential gap in the complexity of proofs in MLK and LK.

Contrary to this intuition, our main result of Chapter 3 proves that the gap is at most quasipolynomial. More precisely, we show that if a sequent of monotone formulas has a proof of size m in LK, then the same sequent has a proof of size $m^{O(\log(m))}$ in MLK. We say that MLK quasipolynomially simulates LK on monotone sequents. Moreover, if only the number of lines is considered, instead of the size, the simulation is then polynomial. A weakness of the result, as stated here, is that the simulation makes sense only on sequents of monotone formulas. To overcome this weakness, we propose to translate each sequent into a set of contradictory clauses (with additional variables), and then use the refutational version of our result which is still true. The reader will observe that each clause may be written as a monotone sequent by putting all negated variables to the left of the sequent (without negation) and all positive variables to the right of the sequent.

Ideally, one would like to prove that MLK polynomially simulates LK on monotone sequents (or on contradictory sets of clauses). Although we are not able to prove this result, we are able to prove an important consequence of it. Namely, that MLK is polynomially bounded if and only if LK and all Frege systems are polynomially bounded. This surprising result indicates that a superpolynomial gap between monotone and non-monotone LK, if any, will be very hard to prove.

The techniques we use to prove these results come from the field of circuit complexity, which is even more surprising since the result goes in the opposite direction there. We use the techniques of slice functions and pseudocomplements. A k-slice function is a Boolean function that is identically 0 on inputs with less than k ones, and identically 1 on inputs with more than k ones. Slice functions have the property that monotone and non-monotone circuits for it have essentially the same size. The reason is that each negated variable can be replaced by a *pseudocomplement* formula which is monotone and behaves properly on inputs with the right number of ones. We adopt these ideas to prove our result.

The main technical step of the proof is showing that certain explicitly given monotone formulas that compute the Boolean threshold functions have short MLK proofs of their basic properties. The k-th Boolean threshold function $\mathrm{TH}_k^n : \{0,1\}^n \rightarrow \{0,1\}$ returns 1 on $x = x_1 \dots x_n \in \{0,1\}^n$ if and only if $\sum_{i=1}^n x_i \geq k$. One of the basic properties with short MLK proofs is that the TH_k^n are symmetric functions. Namely, that for every permutation π on $\{1,\dots,n\}$, we have that $\mathrm{TH}_k^n(x_1,\dots,x_n) = \mathrm{TH}_k^n(x_{\pi(1)},\dots,x_{\pi(n)})$. Since the size of the explicitly given formulas is quasipolynomial $n^{O(\log(n))}$, the whole simulation will be quasipolynomial.

Some other results that we obtain concern the size of proofs of certain specific tautologies. We exhibit a quasipolynomial-size proof of the Pigeonhole Principle sequent PHP_n^{n+1} in MLK. This proof does not use the simulation result and is of independent interest as an alternate to the well-known polynomial upper bound of Buss [Bus97] in Frege systems. In fact, as a byproduct to the generality in which we state our results, we obtain a subexponential-size $2^{O(n^{1/d})}$ proof of PHP_n^{n+1} in depth-d MLK. Although non-monotone upper bounds of this type should easily follow from Buss' techniques, we are not aware of any explicit statement of this even weaker result. Note, on the other hand, that the best lower bound for PHP_n^{n+1} in depth-d Frege systems is of the form $2^{\Omega(n^{1/6^d})}$ [PBI93, KPW95]. We also prove that some specific versions of the Pigeonhole Principle, such as the Onto or the Functional versions, have polynomial-size MLK proofs. We also consider the Matching and Clique-Coclique principles. We show that all these have polynomial-size MLK proofs. As corollaries, we obtain exponential separations between MLK and bounded-depth systems, and MLK and Cutting Planes.

Lower Bounds Beyond Resolution Understanding the expressive power of alternating conjunctions and disjunctions is one of the most important and recurrent themes of Complexity Theory. In Proof Complexity, one considers proof systems whose allowed formulas have a bound on the number of alternations. The extreme case is Resolution which is a depth-0 system. While strong lower bounds are known for proof systems allowing depth-d formulas for small d, the techniques seem unable to produce lower bounds for some important tautologies. The most interesting cases are the Weak Pigeonhole Principle PHP_n^{2n} and Random CNF formulas. These two families of tautologies have been proven hard for Resolution, but resist to be proven hard for depth-d systems for small $d \geq 1$. In the case of PHP_n^{2n}, it should be noted that quasipolynomial-size depth-1 proofs exist, and so "hard" means non-polynomial in this case.

The main result of Chapter 4 breaks the psychological barrier of one alternation for the Weak Pigeonhole Principle and Random CNF formulas with clause density near the threshold. We consider proof systems whose allowed formulas are disjunctions of conjunctions of two literals. A concrete system of this type was introduced by Krajícek [Kra00] under the name Res(2) as an obvious generalization of Resolution. We make an abuse of language and say that a sound proof system with rules of fan-in two whose allowed formulas are of this type is a Res(2)-system. We prove that PHP_n^{2n} requires exponential-size $2^{\Omega(n/(\log(n))^{14})}$ refutations in every Res(2)-system. Similarly, Random CNF formulas with clause density near the threshold ($m = 5n$, say), require refutations of size $2^{\Omega(n^{1/3}/(\log(n))^2)}$ in all Res(2)-systems. We note that these are the strongest lower bounds for the Weak Pigeonhole Principle and Random CNF formulas known to date.

To appreciate the strength of Res(2)-systems, one should examine the behavior of the known techniques to prove lower bounds in depth-d systems. The only known technique for depth-d systems is essentially the random restriction method. It turns out that such a technique, by itself, is unable to produce any lower bound for the Weak Pigeonhole Principle and Random CNF formulas. The reason is that, in order to simplify a formula that involves a single alternation of conjunctions and disjunctions, one needs to apply a relatively large restriction that may trivialize the initial set of clauses. This is better seen in the case of Random CNF formulas. In order to switch a disjunction of conjunctions of two literals into a small conjunction of disjunctions with an exponentially small probability of failure, one needs to apply a restriction of length $\Omega(n)$, and that would falsify an initial 3-clause almost surely. The precise details will be provided in Chapter 4[2].

The way out we suggest is to combine the method of random restrictions with the width method for Resolution. The main observation, whose proof is the main technical contribution of Chapter 4,

[2]For the anxious reader, consider Hastad's Switching Lemma as stated in the Handbook [BS90]. Switching a 2-DNF into an s-CNF requires $p \leq 1/2\gamma \leq 1/10$. This means that each variable is set to 0 or 1 with probability at least $9/20$, which implies a constant probability of falsifying a 3-clause.

is that a disjunction of conjunctions of two literals whose total number of literals is large (whose width is large), is significantly shortened almost surely by a relatively short random restriction. This holds even when restrictions have dependencies among variables, as in the ones used for the Weak Pigeonhole Principle. This is the most difficult and interesting case. In order to prove this claim, we have to resort to the use of martingales and strong concentration bounds such as Azuma's Inequality. To our knowledge, this is the first time this technique is used in the field. Our proof also requires us to show the existence of certain bipartite graphs with strong expansion properties that are robust to a random removal of a small set of nodes and their edges.

Finally, it is known that Res(2) is a system strictly stronger in power than Resolution. Indeed, one can prove a quasipolynomial $m^{\Omega(\log(m)/\sqrt{\log\log(m)})}$ separation [ABE01][3]. A consequence of this result is that our lower bounds would not follow from the known exponential $2^{\Omega(n)}$ lower bounds for the Weak Pigeonhole Principle PHP_n^{2n} in Resolution [BT88] and an alleged efficient simulation of Res(2) in Resolution. Indeed, any such simulation should be quasipolynomial, and that would yield lower bounds of the form $2^{\Omega(n^{\epsilon})}$, instead of the stronger $2^{\Omega(n/(\log(n))^c)}$ that we obtain. We see this observation as an indication that our techniques are really different and stronger than those used for Resolution.

We close Chapter 4 with a result that relates Res(2) to a proof system based on Cutting Planes and the calculus of Lovász-Schrijver introduced by Pudlák [Pud99]. The latter system works with quadratic inequalities. The aim of this result is to shed more light on the observation that the allowed formulas of Res(2) are the Boolean analogue of quadratic inequalities, just as clauses are the Boolean analogue of linear inequalities. To be more precise, we show that the calculus of Lovász-Schrijver augmented with the division rule for quadratic inequalities polynomially simulates Res(2). This should be compared with the known result that Cutting Planes polynomially simulates Resolution. We also introduce a proof system that works with quadratic inequalities but avoids divisions of any type. This system is called Q and we show that it polynomially simulates Res(2) too.

Improved Bounds on WPHP and Infinitely Many Primes We outlined some of the reasons for being interested in the complexity of proofs of the Weak Pigeonhole Principle in Section 1.4. The list is not complete, however. One of the most important challenges in propositional proof complexity is understanding the complexity of proof systems whose allowed formulas have bounded-depth. In a major breakthrough, Ajtai [Ajt88] proved that PHP_n^{n+1} requires superpolynomial-size proofs in bounded-depth systems. The lower bound was significantly strengthened by Pitassi, Beame and Impagliazzo [PBI93] and Krajicek, Pudlák and Woods [KPW95] independently. They proved a lower bound of the form $2^{\Omega(n^{1/6^d})}$ for depth-d systems.

[3]We note that this is in an article co-authored by the author [ABE01], but this result will not be part of this thesis. Actually, a slightly subexponential separation has been proven recently by the author and Bonet [AB01].

For the Weak Pigeonhole Principle PHP_n^{2n}, the situation is quite different. While Buss and Turán [BT88] showed that PHP_n^{2n} requires exponential-size proofs in Resolution, Paris, Wilkie and Woods [PWW88, Kra95] proved that it has quasipolynomial-size $n^{O(\log(n))}$ proofs in bounded-depth systems. More recently, Maciel, Pitassi and Woods [MPW00] gave a new quasipolynomial-size proof of optimal depth. Their system is essentially Res(log), where the allowed formulas are disjunctions of conjunctions of up to $\log(n)^{O(1)}$ literals. Both papers left open, however, whether depth could be traded for size; that is, whether allowing more depth in the proof would allow us reduce the size below $n^{\log(n)}$. We note that such a trade-off is known for the even weaker Pigeonhole Principle $\mathrm{PHP}_n^{n^2}$ [PWW88].

The proofs of Paris, Wilkie and Woods [PWW88] and Maciel, Pitassi and Woods [MPW00] consist in reducing PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$. In both cases, they build an injective map from $\{0, \ldots, n^2 - 1\}$ to $\{0, \ldots, n-1\}$ by repeatedly composing a supposedly injective map from $\{0, \ldots, 2n-1\}$ to $\{0, \ldots, n-1\}$. The difference in their proofs is, essentially, in the proof of $\mathrm{PHP}_n^{n^2}$. Our new contribution is showing that the repeated composition technique can be made more efficient in terms of size. That is, we reduce PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$ in size $n^{o(\log(n))}$ (notice the small oh). The price we need to pay for that is an increase in depth. More precisely, we show that PHP_n^{2n} reduces to $\mathrm{PHP}_n^{n^2}$ in size $n^{O(d(\log(n))^{2/d})}$ and depth d. This gives us the desired size-depth trade-off upper bound for PHP_n^{2n} since $\mathrm{PHP}_n^{n^2}$ is provable in size $n^{O(\log^{(2)}(n))}$ and depth $O(1)$ [PWW88, Kra95].

The most interesting particular case of our size-depth trade-off is when $d = O(1)$ since it proves that the previously known upper bound in bounded-depth Frege is not optimal. Indeed, $n^{O(d(\log(n))^{2/d})}$ grows slower than $n^{c\log(n)}$ for any constants $d > 2$ and $c > 0$. Thus, any lower bound proof will have to focus on a bound weaker than $n^{\log(n)^{\epsilon}}$ for any $\epsilon > 0$. We believe this is valuable information. The other interesting particular case is when $d = O(\log\log(n))$. In that case we obtain a proof of size $n^{O(\log\log(n))}$ and depth $O(\log\log(n))$. The bound $n^{O(\log\log(n))}$ is new in the context of PHP_n^{2n}.

The method that we use to reduce the size of the composition technique is inspired from the theory of automata. We observe that checking whether b is the image of a under repeated composition of a function f is a reachability problem in a graph. Therefore, one can use (an analogue of) Savitch's Theorem to efficiently solve the reachability problem. We are more ambitious and we use ideas from an old theorem of Nepomnjaščij that achieves a size-depth trade-off for the same problem [Nep70]. We formalize the ideas in Nepomnjaščij's Theorem into a theorem of Bounded Arithmetic with an automatic translation into propositional Gentzen Calculus. This formalization may be of independent interest. We note that Nepomnjaščij's Theorem has received a renewed deal of attention recently in the context of time-space trade-off lower bounds for the satisfiability problem [For97, LV99, FvM00].

The new bounds on the Weak Pigeonhole Principle that we obtain have some consequences for Feasible Number Theory whose aim is to develop as much number theory as possible without exponentiation. The main open problem of the field is whether the bounded arithmetic theory $I\Delta_0$ can prove that there are unboundedly many primes [PWW88, MM89, BI91]. The theory $I\Delta_0$ is Peano Arithmetic with induction restricted to Δ_0-formulas, that is, formulas whose quantifiers are of the form $(\forall x \leq y)$ and $(\exists x \leq y)$. Of course, Euclides' proof cannot be carried over in $I\Delta_0$ since exponentially large numbers are required in the proof. In a major breakthrough, Woods [Woo81] showed that exponentiation can be replaced by a combinatorial argument using the Pigeonhole Principle PHP_n^{n+1}, and Paris, Wilkie and Woods [PWW88] realized that the Weak Pigeonhole Principle PHP_n^{2n} was enough for that proof. As a corollary to their results, they show that $I\Delta_0$ augmented with the statement that $x^{\log(x)}$ exists proves that $(\exists y)(y > x \wedge \mathit{prime}(y))$. Our results improve this to show that, for every standard natural number k, the theory $I\Delta_0$ augmented with the statement that $x^{\log(x)^{1/k}}$ exists proves that $(\exists y)(y > x \wedge \mathit{prime}(y))$. Therefore, the large number assumption "$x^{\log(x)}$ exists" is not optimal. Indeed, for $k > 1$, one can build a model of $I\Delta_0$ with a non-standard element a such that $a^{\log(a)^{1/k}}$ exists in the model but $a^{\log(a)}$ does not.

Chapter 2

Preliminaries

2.1 Main Definitions

All propositional proof systems we will deal with can be seen as restrictions of the popular Gentzen Calculus, also called Sequent Calculus in some textbooks. For this reason, we find it useful to define the propositional Gentzen Calculus, denoted LK, in its full generality first.

All our formulas are over the basis $\{\wedge, \vee, \neg\}$ with propositional variables $v_1, v_2, \ldots$ and constants 0 and 1. Formulas may be defined recursively: propositional variables and constants are formulas, if A and B are formulas then $(\neg A)$, $(A \wedge B)$ and $(A \vee B)$ are also formulas, and nothing else is a formula. Alternatively, one may define formulas to be Boolean circuits in tree-form, that is, circuit in which every gate is used at most once. Since we view $\wedge$ and $\vee$ as associative and commutative connectives, we omit unecessary parentheses. Thus, for us, $A \wedge B \wedge C$, $(A \wedge B) \wedge C$, $C \wedge (A \wedge B)$, ... will all be the same formula. We will also omit some parentheses by giving the higher priority to $\neg$. Thus, $\neg A \wedge B$ really means $(\neg A) \wedge B$, and not $\neg(A \wedge B)$.

We may *normalize* Boolean formulas by pushing the negation to the variables using the De Morgan rules, and collapsing consecutive levels of connectives of the same type into a connective of larger *fan-in*, that is, with a larger number of operands. With this approach, the *depth* of a formula is the maximal nesting of conjunctions and disjunctions in a path of the normalized Boolean formula. Thus, literals (propositional variables or negations of propositional variables) have depth 0, conjunctions or disjunctions of literals have depth 1, and so on. The *bottom fan-in* of a formula is the maximum fan-in of a conjunction or disjunction at depth 1.

In the following, Roman uppercase letters such as $A, A', B, \ldots$ represent formulas, and Greek uppercase letters letters such as $\Gamma, \Gamma', \Delta, \ldots$ represent sequences of formulas that we interpret as multisets of formulas.

Axioms:

$$\overline{A \vdash A} \qquad \overline{0 \vdash \Gamma} \qquad \overline{\Gamma \vdash 1}$$

Weakening Rules:

$$\frac{\Gamma \vdash \Delta}{\Gamma, A \vdash \Delta} \qquad \frac{\Gamma \vdash \Delta}{\Gamma \vdash A, \Delta}$$

Cut Rule:

$$\frac{\Gamma \vdash A, \Delta \qquad \Gamma', A \vdash \Delta'}{\Gamma, \Gamma' \vdash \Delta, \Delta'}$$

Left Logical Rules:

$$\frac{\Gamma, A, B \vdash \Delta}{\Gamma, (A \wedge B) \vdash \Delta} \qquad \frac{\Gamma, A \vdash \Delta \qquad \Gamma', B \vdash \Delta'}{\Gamma, \Gamma', (A \vee B) \vdash \Delta, \Delta'} \qquad \frac{\Gamma \vdash A, \Delta}{\Gamma, \neg A \vdash \Delta}$$

Right Logical Rules:

$$\frac{\Gamma \vdash A, B, \Delta}{\Gamma \vdash (A \vee B), \Delta} \qquad \frac{\Gamma \vdash A, \Delta \qquad \Gamma' \vdash B, \Delta'}{\Gamma, \Gamma' \vdash (A \wedge B), \Delta, \Delta'} \qquad \frac{\Gamma, A \vdash \Delta}{\Gamma \vdash \neg A, \Delta}$$

An *LK-proof* of $\Gamma \vdash \Delta$ is a sequence of lines ending with $\Gamma \vdash \Delta$, each of which is either an axiom of LK, or has been obtained by a rule of LK from previous lines in the sequence. Each line of the proof is called a *sequent*. An *LK-refutation* of $\{\Gamma_i \vdash \Delta_i : i \in I\}$ is an LK-proof of the empty sequent $\vdash$ when each $\Gamma_i \vdash \Delta_i$ may be used as an axiom. When we restrict the proof in such a way that each derived sequent can be used only once as a premise in a rule, we say that the proof is *tree-like*. The same definition applies to refutations. The *size* of the proof or refutation is the overall number of *symbols* used in it. In this definition, each occurrence of a propositional variable counts one, and each connective also counts one. Parentheses do not count in the size. The *depth* of a proof is the maximum depth of a formula appearing in it.

2.2 Resolution

One of the most popular proof systems is Resolution. This system works with *clauses*, that is, disjunctions of variables or negated variables which are also called *literals*. The goal is to refute a given set of clauses by deriving the empty clause. There is a unique rule of inference, the so-called Resolution rule:

Resolution Rule:

$$\frac{A \vee x \qquad B \vee \neg x}{A \vee B},$$

where A and B are clauses, and x is a propositional variable. As in LK, a Resolution refutation is tree-like if every clause in the refutation is used at most once as the premise of a rule. The length of a Resolution refutation is the number of clauses in it, and the size is the total number of symbols in it.

Resolution may be seen as the restriction of LK when all formulas in the sequents are variables. For a given clause C of the form $\neg x_1 \vee \ldots \vee \neg x_r \vee x_{r+1} \vee \ldots \vee x_s$, let $s(C)$ be the sequent $x_1, \ldots, x_r \vdash x_{r+1}, \ldots, x_s$. Similarly for a given sequent S of the form $x_1, \ldots, x_r \vdash x_{r+1}, \ldots, x_s$, let $c(S)$ be the clause $\neg x_1 \vee \ldots \vee \neg x_r \vee x_{r+1} \vee \ldots \vee x_s$.

Proposition 1 *Let $\mathcal{C} = \{C_1, \ldots, C_m\}$ be a set of clauses and let $\mathcal{S} = \{S_1, \ldots, S_m\}$ be a set of sequents such that $s(C_i) = S_i$ and $c(S_i) = C_i$. If $\mathcal{C}$ has a Resolution refutation of size S, then $\mathcal{S}$ has an LK refutation of size $O(S)$ in which every formula is a propositional variable. Conversely, if $\mathcal{S}$ has an LK refutation of size S in which every formula is a propositional variable, then $\mathcal{C}$ has a Resolution refutation of size $O(S)$.*

Proof: The first implication is very easy: replace each clause in the Resolution refutation by its translation into a sequent, and replace each application of the Resolution rule by a cut.

For the second implication, first replace each sequent by its translation into a clause. Observe that the logical rules of introduction of $\neg$, $\wedge$ and $\vee$ cannot be used in the LK proof since all formulas there are propositional variables. Similarly, the axioms about the constants 0 and 1 cannot be used since these are not variables. Thus, all axioms are of the form $x \vdash x$ for some variable x, and the only rules that are left are the weakening rules and cut. By induction on the length of the proof, one can easily show how to eliminate each weakening rule by standard proof-theoretic techniques. Similarly, one can ignore all cuts with an axiom $x \vdash x$ since that would not change the sequent. The resulting sequence of clauses is a Resolution refutation. $\square$

In view of this Proposition, the restriction of LK so that the only allowed formulas in the sequents are propositional variables is equivalent to Resolution for all practical purposes.

2.3 Extensions of Resolution.

Krajíček introduced a refutational proof system called Res(k) that works with k-disjunctions, that is, disjunctions of conjunctions of up to k literals. The goal, again, is to derive the empty clause from a given set of clauses. There are three inference rules in Res(k): Weakening, Introduction of $\wedge$ and Cut.

Weakening:

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

Introduction of $\wedge$:

$$\frac{A \vee (l_1 \wedge \ldots \wedge l_h) \qquad B \vee (l_{h+1} \wedge \ldots \wedge l_r)}{A \vee B \vee (l_1 \wedge \ldots \wedge l_r)}$$

Cut:
$$\frac{A \vee (l_1 \wedge \ldots \wedge l_r) \qquad B \vee \neg l_1 \vee \ldots \vee \neg l_r}{A \vee B},$$
where A and B are k-disjunctions, $1 \leq h < r \leq k$, and the l_i's are literals. We also allow axioms of the form $x \vee \neg x$. The length of a Res(k) refutation is the number of k-disjunctions in it, and the size is the number of symbols in it.

$\mathrm{Res}(k)$ may be seen as the restriction of LK so that all formulas in the sequents are depth-1 formulas with fan-in at most k. That is, all formulas in the sequents are either conjunctions of up to k literals, or disjunctions of up to k literals. Given a k-disjunction D of the form $\bigvee_{i=1}^{t} \bigwedge_{j \in L_i} l_{i,j}$, let $s(D)$ be the sequent $\vdash \bigvee_{i=1}^{t} \bigwedge_{j \in L_i} l_{i,j}$. Similarly, given a sequent S of the form $C_1, \ldots, C_r \vdash C_{r+1}, \ldots, C_s$ where each C_i is a depth-1 formula of fan-in at most k, let $d(S)$ be the k-disjunction $\bigvee_{i=1}^{r} \neg C_i \vee \bigvee_{i=r+1}^{s} C_i$, where $\neg C$ stands for $\neg l_1 \vee \ldots \vee \neg l_t$ if $C = l_1 \wedge \ldots \wedge l_t$, and for $\neg l_1 \wedge \ldots \wedge \neg l_t$ if $C = l_1 \vee \ldots \vee l_t$.

Proposition 2 *Let $\mathcal{C} = \{C_1, \ldots, C_m\}$ be a set of clauses and let $\mathcal{S} = \{S_1, \ldots, S_m\}$ be a set of sequents so that $s(C_i) = S_i$ and $d(S_i) = C_i$. If $\mathcal{C}$ has a Res(k) refutation of size S, then $\mathcal{S}$ has an LK refutation of size $O(S)$ in which every formula in the sequent is a depth-1 formula of fan-in at most k. Conversely, if $\mathcal{S}$ has an LK refutation of size S in which every formula is a depth-1 formula of fan-in at most k, then $\mathcal{C}$ has a Res(k) refutation of size $O(S)$.*

Proof: The first implication is almost immediate. First, translate each k-disjunction into a sequent. Each application of the weakening rule gets translated into an application of weakening and a right introduction of $\vee$. Each application of introduction of $\wedge$ gets translated into an application of right introduction of $\wedge$ followed by right introduction of $\vee$. Do the same for each cut. Finally, the axioms $x \vee \neg x$ become $\vdash x \vee \neg x$, and these have pretty short proofs in LK.

For the second implication, first replace each sequent by its translation into a k-disjunction. In the LK proof, the rule of introduction of $\neg$ was only used over literals since all formulas must be depth-1 formulas of fan-in at most k. By the translation, these rules may be ignored. The rule of right introduction of $\vee$ may be simulated by a weakening. Similarly, the rule of left introduction of $\wedge$ may be simulated by a weakening since left conjunctions turn into disjunction by the translation. The rule of right introduction of $\wedge$ may be simulated by an introduction of $\wedge$. Similarly, the rule of left introduction of $\vee$ may be simulated by an introduction of $\wedge$ since left disjunctions turn into conjunctions by the translation. Cuts may very well be translated by cuts, and weakening rules may be simulated or ignored. Finally, we need to simulate an axiom of the form $A \vdash A$. Suppose A is $l_1 \wedge \ldots \wedge l_r$. The sequent $A \vdash A$ becomes $\neg l_1 \vee \ldots \vee \neg l_r \vee (l_1 \wedge \ldots \wedge l_r)$. Start with $\neg l_i \vee l_i$ and introduce r conjunctions. This completes the simulation. $\square$

As it is the case in Resolution, the refutation system Res(k) and LK in which all formulas have depth 1 and fan-in at most k are equivalent for all practical purposes.

2.4 Bounded-Depth Gentzen Calculus.

Depth-k LK is the proof system that results from LK when we restrict all formulas in the sequents to have depth at most k. The corresponding refutation system is also called Depth-k LK. It will be explicitly said whether we are talking about proofs or refutations.

We note that refutational Depth-0 LK is equivalent to Resolution. The reason is that Depth-0 LK is equivalent to the restriction of LK to sequents whose formulas are propositional variables only. To see this, simply ignore introductions of negation, and invert some weakenings if necessary. On the other hand, Depth-1 LK is equivalent to Res(n) when n is the total number of propositional variables in the formulas. Some authors use the notation Depth-$(d+0.5)$ LK to mean the restriction of LK to sequents whose formulas have depth bounded by $d+1$ and bottom fan-in bounded by a polylogarithm of the number of variables. The system refutational Depth-0.5 LK is particularly interesting since it coincides with Res($\log$) in the notation of Krajícek [Kra00]. There, Res($\log$) is meant to be $\bigcup_{k\geq 0} \mathrm{Res}((\log n)^k)$ where n is the number of variables in the formulas.

2.5 Monotone Gentzen Calculus and Intuitionistic LK.

Pudlák [Pud99] introduced a monotone version of the propositional calculus. He suggested to restrict LK to monotone formulas only, that is, to formulas that do not use negation. The resulting proof system was called MLK in [AGG01]. Observe that Resolution is equivalent to a subsystem of refutational MLK. Indeed, if only propositional variables are allowed as formulas, these are necessarily monotone.

The Intuitionistic Calculus JK is the proof system that results from LK when we restrict all sequents to have at most one formula on its right-hand side. For this system to make good sense, one needs to redefine the rule of Right Weakening. The new rule is as follows:

Right Weakening Rule of JK:

$$\frac{\Gamma \vdash F}{\Gamma \vdash (A \vee F)}$$

It is known that JK is interpretable in LK in the sense that there exists a translation of propositional formulas $A \mapsto t(A)$ so that if $\vdash A$ has an LK-proof, then $\vdash t(A)$ has a JK-proof (see [Tak87], for example). Moreover, the sizes of the proofs are polynomially related. On the other hand, every MLK-proof of a monotone sequent whose right-hand side consists of a single formula may be translated into a JK-proof of roughly the same size. This result is due to Bílková [B0́1]. Thus, MLK is a subsystem of JK.

Chapter 3

The Monotone Calculus

3.1 Threshold Formulas

The goal of this section is twofold. First, we define some explicit monotone formulas for the Boolean threshold functions with a nice size-depth trade-off property. Second, we show that these formulas admit reasonably short monotone proofs of their basic properties. In particular, we prove that they define symmetric functions.

For every n and $k \in \{0, \ldots, n\}$, let $\mathrm{TH}^n_k : \{0,1\}^n \to \{0,1\}$ be the Boolean function that returns 1 on $x = x_1 \ldots x_n \in \{0,1\}^n$ if and only if $\sum_{i=1}^k x_i \geq k$. Each TH^n_k is called a threshold function. Valiant [Val84] proved that every threshold function TH^n_k is computable by a monotone formula of size polynomial in n. The proof is probabilistic, and so the construction is not explicit. In the same paper, Valiant mentioned that a divide and conquer strategy leads to explicit quasipolynomial-size monotone formulas for all threshold functions. The same construction appears in the book by Wegener [Weg87], and in the more recent book by Vollmer [Vol99]. Here we revisit that construction with a minor modification to achieve a size-depth trade-off.

Theorem 1 *Let $s, r \in \mathbb{N}$ and $n = s^r$. For every $k \in \{0, \ldots, n\}$, there exist monotone formulas $\mathrm{th}^n_{k,s}(x_1, \ldots, x_n)$ computing $\mathrm{TH}^n_k(x_1, \ldots, x_n)$ of size $n^{s \log_s(n)}$ and depth $2\log_s(n)$.*

Proof: In the following, we use the notation $\overline{i}$ to denote a sequence $(i_1, \ldots, i_l)$. The length l of the sequence will be clear from context. Define $\mathrm{th}^1_{0,s}(x) = 1$ and $\mathrm{th}^1_{1,s}(x) = x$. For every $k \leq n$, we let $\mathrm{th}^n_{k,s}(x_1, \ldots, x_n)$ be the formula

$$\bigvee_{\overline{i} \in I^n_{k,s}} \bigwedge_{j=1}^{s} \mathrm{th}^{n/s}_{i_j,s}(x_{(j-1)n/s+1}, \ldots, x_{jn/s}),$$

where $I^n_{k,s} = \{(i_1, \ldots, i_s) \in \mathbb{N}^s : 0 \leq i_j \leq n/s, \ \sum_j i_j \geq k\}$. It is straightforward to prove, by induction on n, that the formula $\mathrm{th}^n_{k,s}(x_1, \ldots, x_n)$ computes the boolean function TH^n_k. The depth of $\mathrm{th}^n_{k,s}(x_1,\ldots,x_n)$ is bounded by $2\log_s(n)$. Moreover, the maximum size of $\mathrm{th}^n_{k,s}(x_1,\ldots,x_n)$, say $S(n,s)$, satisfies the recurrence

$$S(n,s) \quad \leq \quad n^s S(n/s,s),$$

so we have $S(n,s) \leq n^{s\log_s(n)}$. □

Observe that when $s = 2$, the depth is bounded by $2\log_2(n)$ and the size is bounded by $n^{2\log_2(n)}$, and that when $s = n^{2/d}$ for a constant $d > 2$, the depth is bounded by d and the size is bounded $2^{n^{3/d}}$.

We establish a number of lemmas stating that the elementary properties of the threshold formulas admit short MLK-proofs. The first property is about monotone formulas in general.

Let A and $B_1,\ldots,B_n$ be formulas, and let $x_1,\ldots,x_n$ be propositional variables that may or may not occur in A. We let $A(x_1/B_1,\ldots,x_n/B_n)$ denote the formula that results from A when all occurrences of x_i (if any) are replaced by B_i (replacements are made simultaneously). Observe that if A and B are monotone formulas, then $A(x/B)$ is also monotone. The non-monotone version of the following lemma appears in [BDG$^+$99, Bus95] (monotonicity is only needed in part (v)).

Lemma 1 *If A is a monotone formula, the sequents (i) $A, x \vdash A(x/1)$, (ii) $A \vdash x, A(x/0)$, (iii) $A(x/1), x \vdash A$, (iv) $A(x/0) \vdash x, A$, and (v) $A(x/0) \vdash A(x/1)$, have MLK-proofs of size quadratic in the size of A and the same depth as A.*

Proof: We will only prove (i) and (v). The rest are similar. If A does not contain the variable x, there is nothing to prove since then $A = A(x/0) = A(x/1)$. Suppose then that x appears in A. We proceed by induction on the construction of A. If A is the propositional variable x, then $A(x/1) = 1$ and so (i) and (v) are simply axioms. Suppose next that $A = B \wedge C$. By induction hypothesis, we have $B, x \vdash B(x/1)$ and $C, x \vdash C(x/1)$. Right $\wedge$-introduction gives $B, C, x \vdash B(x/1) \wedge C(x/1)$, and left $\wedge$-introduction gives the result. Similarly, $B(x/0) \vdash B(x/1)$ and $C(x/0) \vdash C(x/1)$ by induction hypothesis. So right $\wedge$-introduction and left $\wedge$-introduction gives the desired result. The case $A = B \vee C$ is no less trivial. □

To simplify notation, in this and following sections we omit the subscript s in proofs, as it is always the same. The first properties are easy:

Lemma 2 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be an exact power of s. Let $0 \leq i_1,\ldots,i_s \leq n/s$, let $k = \sum_j i_j$, and let $h, l \in \mathbb{N}$ with $n \geq h \geq l$. The sequents*

(i) $\vdash \mathrm{th}^n_{0,s}(x_1,\ldots,x_n)$,

(ii) $\mathrm{th}^n_{n,s}(x_1,\ldots,x_n) \vdash \bigwedge_{i=1}^n x_i$,

(iii) $\bigwedge_{j=1}^{s} \mathrm{th}_{i_j,s}^{n/s}(x_{(j-1)n/s+1}, \ldots, x_{jn/s}) \vdash \mathrm{th}_{k,s}^{n}(x_1, \ldots, x_n)$,

(iv) $\mathrm{th}_{h,s}^{n}(x_1, \ldots, x_n) \vdash \mathrm{th}_{l,s}^{n}(x_1, \ldots, x_n)$,

have MLK-proofs of size polynomial in $n^{s \log_s(n)}$ *and depth* $2\log_s(n)$*.*

In the next lemmas we give MLK-proofs of the basic properties relative to the symmetry of the threshold formulas (Theorem 2 below).

Lemma 3 *Let* $s \in \mathbb{N}$ *and let* $n \in \mathbb{N}$ *be an exact power of* s*. Let* $m, k, l \in \mathbb{N}$ *with* $0 < m \leq n$*,* $0 \leq k < n$*, and* $1 \leq l \leq n$*. The sequents*

(i) $\mathrm{th}_{k+1,s}^{n}(x_1, \ldots, x_l/1, \ldots, x_n) \vdash \mathrm{th}_{k,s}^{n}(x_1, \ldots, x_l/0, \ldots, x_n)$

(ii) $\mathrm{th}_{m-1,s}^{n}(x_1, \ldots, x_l/0, \ldots, x_n) \vdash \mathrm{th}_{m,s}^{n}(x_1, \ldots, x_l/1, \ldots, x_n)$

have MLK-proofs of size polynomial in $n^{s \log_s(n)}$ *and depth* $2\log_s(n)$*.*

Proof: In the following, let $\overline{x}_j = (x_{(j-1)n/s+1}, \ldots, x_{jn/s})$ for every $j \in \{1, \ldots, s\}$, and let $\overline{x}'_j$ be the result of replacing x_l by 0 in $\overline{x}_j$. We first show (i). We use induction on n, where the base case is $\mathrm{th}_1^1(1) \vdash \mathrm{th}_0^1(0)$. Assume without loss of generality that $l \leq n/s$, that is, x_l is in the first block of variables $\overline{x}_1$. Recall the definition of $\mathrm{th}_{k+1}^{n}(x_1, \ldots, x_n)$:

$$\bigvee_{\overline{i} \in I_{k+1}^{n}} \bigwedge_{j=1}^{s} \mathrm{th}_{i_j}^{n/s}(\overline{x}_j).$$

Fix $(i_1, \ldots, i_s) \in I_{k+1}^{n}$. If $i_1 = 0$, then $\sum_{j=2}^{s} i_j \geq k+1$ so that $i_q > 0$ for some $q \in \{2, \ldots, s\}$. Then, $\mathrm{th}_{i_q}^{n/s}(\overline{x}_q) \vdash \mathrm{th}_{i_q-1}^{n/s}(\overline{x}_q)$ by part (iv) of Lemma 2. On the other hand, clearly $\mathrm{th}_{i_j}^{n/s}(\overline{x}_j) \vdash \mathrm{th}_{i_j}^{n/s}(\overline{x}_j)$ for every $j \in \{2, \ldots, s\} - \{q\}$. Moreover, we have $\vdash \mathrm{th}_0^{n/s}(\overline{x}'_1)$ by part (i) of Lemma 2. Note, by the way, that $\overline{x}'_j = \overline{x}_j$ for every $j \in \{2, \ldots, s\}$. Right $\wedge$-introduction, left weakening, and left $\wedge$-introduction gives then

$$\bigwedge_{j=1}^{s} \mathrm{th}_{i_j}^{n/s}(\overline{x}'_j) \vdash \mathrm{th}_0^{n/s}(\overline{x}'_1) \wedge \bigwedge_{2 \leq j < q} \mathrm{th}_{i_j}^{n/s}(\overline{x}'_j) \wedge \mathrm{th}_{i_q-1}^{n/s}(\overline{x}'_q) \wedge \bigwedge_{q < j \leq s} \mathrm{th}_{i_j}^{n/s}(\overline{x}'_j).$$

A cut with part (iii) of Lemma 2 gives $\bigwedge_{j=1}^{s} \mathrm{th}_{i_j}^{n/s}(\overline{x}'_j) \vdash \mathrm{th}_{t-1}^{n}(\overline{x}'_1, \ldots, \overline{x}'_s)$, where $t = \sum_{j=1}^{s} i_j$. Finally, since $t - 1 \geq k + 1 - 1 = k$, a cut with part (iv) of Lemma 2 gives the result.

If $i_1 > 0$, we use the induction hypothesis on n to get $\mathrm{th}_i^{n/s}(\overline{x}'_1) \vdash \mathrm{th}_{i-1}^{n/s}(\overline{x}'_1)$. Easy manipulation as before gives

$$\bigwedge_{j=1}^{s} \mathrm{th}_{i_j}^{n/s}(\overline{x}'_j) \vdash \mathrm{th}_{i_1-1}^{n/s}(\overline{x}'_1) \wedge \bigwedge_{2 \leq j \leq s} \mathrm{th}_{i_j}^{n/s}(\overline{x}'_j).$$

Finally an application of parts (iii) and (iv) of Lemma 2 gives the desired result. The proof of (ii) is very similar. □

Lemma 4 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be an exact power of s. Let $m, k, l \in \mathbb{N}$ with $1 \leq k < l \leq n$, and $m \leq n$, the sequents*

(i) $\mathrm{th}^n_{m,s}(x_1, \ldots, x_k/1, \ldots, x_l/0, \ldots, x_n) \vdash \mathrm{th}^n_{m,s}(x_1, \ldots, x_k/0, \ldots, x_l/1, \ldots, x_n)$

(ii) $\mathrm{th}^n_{m,s}(x_1, \ldots, x_k/0, \ldots, x_l/1, \ldots, x_n) \vdash \mathrm{th}^n_{m,s}(x_1, \ldots, x_k/1, \ldots, x_l/0, \ldots, x_n)$

have MLK-proofs of size polynomial in $n^{s \log_s(n)}$ and depth $2\log_s(n)$.

Proof: We use the same notation as in the proof of Lemma 3. Both proofs are identical. It is enough to prove (i) when x_k and x_l fall in different blocks of variables. The complete proof of (i) would then be a simple induction on the recursive definition of $\mathrm{th}^n_m(x_1, \ldots, x_n)$ whose base case is when that happens. Notice that the base case is eventually reached, at latest, when $n = s$. So assume x_k and x_l fall in blocks a and b respectively. In the following, let $\overline{x}_a/1$ be the result of replacing x_k by 1 in $\overline{x}_a$, and define the notation $\overline{x}_a/0$, $\overline{x}_b/1$ and $\overline{x}_b/0$ analogously. Recall the definition of $\mathrm{th}^n_m(\overline{x}_1, \ldots, \overline{x}_s)$:

$$\bigvee_{\overline{i} \in I^n_m} \bigwedge_{j=1}^{s} \mathrm{th}^{n/s}_{i_j}(\overline{x}_j).$$

Fix $(i_1, \ldots, i_s) \in I^n_m$. If $i_a > 0$, then Lemma 3 shows that $\mathrm{th}^{n/s}_{i_a}(\overline{x}_a/1) \vdash \mathrm{th}^{n/s}_{i_a-1}(\overline{x}_a/0)$. Similarly, whenever $i_b < n/s$ we have $\mathrm{th}^{n/s}_{i_b}(\overline{x}_b/0) \vdash \mathrm{th}^{n/s}_{i_b+1}(\overline{x}_b/1)$. From these two sequents, the result follows easily when $i_a > 0$ and $i_b < n/s$. Consider next the case in which either $i_a = 0$ or $i_b = n/s$. If $i_b = n/s$, then $\mathrm{th}^{n/s}_{i_b}(\overline{x}_b/0)$ is just provably false by part (ii) of Lemma 2, and the result follows easily. If $i_a = 0$, then $\mathrm{th}^{n/s}_{i_a}(\overline{x}_{i_a}/1)$ is just provably true by part (i) of Lemma 2. On the other hand, $\mathrm{th}^{n/s}_{i_b}(\overline{x}_b/0) \vdash \mathrm{th}^{n/s}_{i_b}(\overline{x}_b/1)$ follows by part (v) of Lemma 1, and the result follows too. $\square$

Lemma 5 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be an exact power of s. Let $m, i, j \in \mathbb{N}$, with $m \leq n$ and $1 \leq i < j \leq n$. The sequent*

$$\mathrm{th}^n_{m,s}(x_1, \ldots, x_i, \ldots, x_j, \ldots, x_n) \vdash \mathrm{th}^n_{m,s}(x_1, \ldots, x_j, \ldots, x_i, \ldots, x_n)$$

has an MLK-proof of size polynomial in $n^{s \log_s(n)}$ and depth $2\log_s(n)$.

Proof: We split the property according to the four possible truth values of x_i and x_j. Namely, we will give proofs of the following four sequents from which the lemma is immediately obtained by the cut rule.

(i) $\mathrm{th}^n_{m,s}(x_1, \ldots, x_i, \ldots, x_j, \ldots, x_n), x_i, x_j \vdash \mathrm{th}^n_{m,s}(x_1, \ldots, x_j, \ldots, x_i, \ldots, x_n)$,

(ii) $\mathrm{th}^n_{m,s}(x_1, \ldots, x_i, \ldots, x_j, \ldots, x_n), x_i \vdash x_j, \mathrm{th}^n_{m,s}(x_1, \ldots, x_j, \ldots, x_i, \ldots, x_n)$,

(iii) $\mathrm{th}^n_{m,s}(x_1, \ldots, x_i, \ldots, x_j, \ldots, x_n), x_j \vdash x_i, \mathrm{th}^n_{m,s}(x_1, \ldots, x_j, \ldots, x_i, \ldots, x_n)$,

(iv) $\mathrm{th}^n_{m,s}(x_1,\ldots,x_i,\ldots,x_j,\ldots,x_n) \vdash x_i, x_j, \mathrm{th}^n_{m,s}(x_1,\ldots,x_j,\ldots,x_i,\ldots,x_n)$.

We only show sequent (ii), the rest of sequents have similar proofs. Two applications of Lemma 1 give

$$\mathrm{th}^n_m(x_1,\ldots,x_i,\ldots,x_j,\ldots,x_n), x_i \vdash x_j, \mathrm{th}^n_m(x_1,\ldots,1,\ldots,0,\ldots,x_n).$$

Lemma 4 gives

$$\mathrm{th}^n_m(x_1,\ldots,x_i,\ldots,x_j,\ldots,x_n), x_i \vdash x_j, \mathrm{th}^n_m(x_1,\ldots,0,\ldots,1,\ldots,x_n).$$

Two more applications of Lemma 1 give

$$\mathrm{th}^n_m(x_1,\ldots,0,\ldots,1,\ldots,x_n), x_i \vdash x_j, \mathrm{th}^n_m(x_1,\ldots,x_j,\ldots,x_i,\ldots,x_n).$$

Finally, a cut between the last two sequents gives (ii). $\Box$

Since every permutation on $\{1,\ldots,n\}$ can be obtained as the composition of (polynomially many) permutations in which only two elements are permuted (transpositions), Lemma 5 easily implies the following theorem.

Theorem 2 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be an exact power of s. Let $m \in \mathbb{N}$, with $m \leq n$, and let π be a permutation over $\{1,\ldots,n\}$. The sequent*

$$\mathrm{th}^n_{m,s}(x_1,\ldots,x_n) \vdash \mathrm{th}^n_{m,s}(x_{\pi(1)},\ldots,x_{\pi(n)})$$

has an MLK-proof of size polynomial in $n^{s\log_s(n)}$ and depth $2\log_s(n)$.

The next two properties state that the smallest threshold formulas are provably equivalent to their usual formulas.

Lemma 6 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be an exact power of s. The sequents*

(i) $\bigvee_i x_i \dashv\vdash \mathrm{th}^n_{1,s}(x_1,\ldots,x_n)$;

(ii) $\bigvee_{i\neq j}(x_i \wedge x_j) \dashv\vdash \mathrm{th}^n_{2,s}(x_1,\ldots,x_n)$;

have MLK-proofs of size polynomial in $n^{s\log_s(n)}$ and depth $2\log_s(n)$.

Proof: All proofs are by induction on n. For (i.$\vdash$), reason as follows. Clearly, $x_1 \vdash \mathrm{th}^1_1(x_1)$ so that the base case holds. Assume then $n > 1$, and that the claim holds for smaller n. Fix

$j_0 \in \{1, \ldots, s\}$. Since $\vdash \mathrm{th}_0^{n/s}(x_{(j-1)n/s+1}, \ldots, x_{jn/s})$ by part (i) of Lemma 2, right $\wedge$-introduction on the induction hypothesis for $\mathrm{th}_1^{n/s}(x_{(j_0-1)n/s+1}, \ldots, x_{j_0 n/s})$ gives

$$\bigvee_{i=1}^{n/s} x_{(j_0-1)n/s+i} \vdash \bigwedge_{\substack{j=1 \\ j \neq j_0}}^{s} \mathrm{th}_0^{n/s}(x_{(j-1)n/s+1}, \ldots, x_{jn/s}) \wedge \mathrm{th}_1^{n/s}(x_{(j_0-1)n/s+1}, \ldots, x_{j_0 n/s}).$$

A cut with part (iii) of Lemma 2 gives $\mathrm{th}_1^n(x_1, \ldots, x_n)$ on the right. Left $\vee$-introduction for each $j_0 \in \{1, \ldots, s\}$ gives then (i). The proof of (i.$\dashv$) is also by induction on n. In fact, we prove the slightly stronger statement: $\mathrm{th}_k^n(x_1, \ldots, x_n) \vdash \bigvee_i x_i$ for every $k \in \{1, \ldots, n\}$. Fix $(i_1, \ldots, i_s) \in I_1^n$, so that $i_j \geq 1$ for some $j \in \{1, \ldots, s\}$. Then, by induction hypothesis, $\mathrm{th}_{i_j}^{n/s}(x_{(j-1)n/s+1}, \ldots, x_{jn/s}) \vdash \bigvee_{i=1}^s x_{(j-1)n/s+i}$. Left weakening, left $\wedge$-introduction, right weakening, and right $\vee$-introduction gives

$$\bigwedge_{j=1}^{s} \mathrm{th}_{i_j}^{n/s}(x_{(j-1)n/s+1}, \ldots, x_{jn/s}) \vdash \bigvee_{i=1}^{n} x_i.$$

Since this was true for an arbitrary $(i_1, \ldots, i_s) \in I_1^n$, the result follows by left $\vee$-introduction. The proof of (ii) is similar and relies on part (i). $\Box$

The next lemma states that threshold functions split by cases:

Lemma 7 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be an exact power of s. Let $m \in \mathbb{N}$ be an exact multiple of s with $m \leq n$. The sequents*

(i) $\mathrm{th}_{m+1,s}^n(x_1, \ldots, x_n) \vdash \mathrm{th}_{m/s+1,s}^{n/s}(\overline{x}_1), \ldots, \mathrm{th}_{m/s+1,s}^{n/s}(\overline{x}_s)$,

(ii) $\mathrm{th}_{m,s}^n(x_1, \ldots, x_n) \vdash \mathrm{th}_{m/s+1,s}^{n/s}(\overline{x}_1), \ldots, \mathrm{th}_{m/s+1,s}^{n/s}(\overline{x}_{s-1}), \mathrm{th}_{m/s,s}^{n/s}(\overline{x}_s)$,

where $\overline{x}_j = (x_{(j-1)n/s+1}, \ldots, x_{jn/s})$, have MLK-proofs of size polynomial in $n^{s \log_s(n)}$ and depth $2\log_s(n)$.

Proof: We first prove (i). Fix $(i_1, \ldots, i_s) \in I_{m+1}^n$. Since m is an exact multiple of s, there must exist a $q \in \{1, \ldots, s\}$ such that $i_q \geq m/s + 1$ for otherwise $\sum_{j=1}^s i_j \leq m$. Then, $\mathrm{th}_{i_q}^{n/s}(\overline{x}_q) \vdash \mathrm{th}_{m/s+1}^{n/s}(\overline{x}_q)$ by part (iv) of Lemma 2. The sequent

$$\bigwedge_{j=1}^{s} \mathrm{th}_{i_j}^{n/s}(\overline{x}_j) \vdash \mathrm{th}_{m/s+1}^{n/s}(\overline{x}_1), \ldots, \mathrm{th}_{m/s+1}^{n/s}(\overline{x}_s)$$

follows by right weakening, left weakening, and left $\wedge$-introduction. Since this happens for every $\overline{i} \in I_{m+1}^n$, the result follows by left $\vee$-introduction. The proof of (ii) is extremely similar. Given $(i_1, \ldots, i_s) \in I_m^n$, either $i_q \geq m/s + 1$ for some $q \in \{1, \ldots, s-1\}$, or $i_s \geq m/s$ for otherwise $\sum_{j=1}^s i_j < m$. Manipulation as in part (i) gives property (ii). $\Box$

3.2 The Pigeonhole Principle

In this section we exhibit a monotone proof of the Pigeonhole Principle using the explicit monotone threshold formulas, and their properties. The Pigeonhole Principle states that if $n+1$ pigeons go into n holes, then there is some hole with more than one pigeon. If $p_{i,j}$ is a propositional variable meaning that pigeon i sits in hole j, the principle is encoded by the following monotone sequent:

$$\bigwedge_{i=1}^{n+1} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i\neq j}}^{n+1} (p_{i,k} \wedge p_{j,k}).$$

From now on we refer to the left part of the sequent as LPHP, and to the right part of the sequent as RPHP. The sequent itself is denoted PHP_n^{n+1}. We need a technical lemma saying that PHP_n^{n+1} can be reduced to the case in which n is an exact power of s.

Lemma 8 *There exists a polynomial $p(n)$ such that, for every $m, S \in \mathbb{N}$, if the sequent PHP_m^{m+1} has an MLK-proof of size S, then, for every $n \leq m$, the sequent PHP_n^{n+1} has an MLK-proof of size at most $S + p(n)$.*

Proof: Suppose that there is a monotone proof $\Psi_1, \Psi_2, \ldots, \mathrm{PHP}_m^{m+1}$ of size at most S, where each Ψ_i is a monotone sequent $\Sigma_i \vdash \Gamma_i$. We get a proof of PHP_n^{n+1} from the proof of PHP_m^{m+1} by replacing some variables by constants as follows. Define a partial truth assignment σ as indicated next. Let $\sigma(p_{k+1,k}) = 1$ for every $k \in \{n+1, \ldots, m\}$. Similarly, for every $k \in \{n+2, \ldots, m+1\}$ and $i \in \{1, \ldots, k-2\}$, let $\sigma(p_{k,i}) = 0$; and for every $i \in \{n+1, \ldots, m\}$ and $k \in \{1, \ldots, i\}$, let $\sigma(p_{k,i}) = 0$. Any other variable remains undefined by σ. Given a sequent $\Sigma \vdash \Gamma$, let $[\Sigma \vdash \Gamma][\sigma]$ be the result of replacing each occurrence of the variable $x \in \mathrm{Dom}(\sigma)$ in Σ or Γ by $\sigma(x)$. The sequence $[\Sigma_1 \vdash \Gamma_1][\sigma], [\Sigma_2 \vdash \Gamma_2][\sigma], \ldots, [\mathrm{PHP}_m^{m+1}][\sigma]$ is a valid proof of $[\mathrm{PHP}_m^{m+1}][\sigma]$. To see this, observe that the initial axioms of the form $p_{i,j} \vdash p_{i,j}$ become $0 \vdash 0$, $1 \vdash 1$, or stay $p_{i,j} \vdash p_{i,j}$, which are all true sequents. Moreover, it is not difficult to give a proof of

$$\left[\bigwedge_{i=1}^{n+1} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigwedge_{i=1}^{m+1} \bigvee_{j=1}^{m} p_{i,j}\right][\sigma]$$

and

$$\left[\bigvee_{k=1}^{m} \bigvee_{\substack{i,j=1 \\ j\neq i}}^{m+1} (p_{i,k} \wedge p_{j,k}) \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ j\neq i}}^{n+1} (p_{i,k} \wedge p_{j,k})\right][\sigma]$$

from the axioms $0 \vdash$ and $\vdash 1$. For example, $[\vdash \bigvee_{j=1}^{m} p_{n+2,j}][\sigma]$ is derivable since $\sigma(p_{n+2,n+1}) = 1$. Two cuts give a proof of PHP_n^{n+1} of size at most $S + p(n)$ for some polynomial $p(n)$, as desired. $\square$

Theorem 3 *Let $s \in \mathbb{N}$ and let $n \in \mathbb{N}$ be such that $s \leq n$. The sequents PHP_n^{n+1} have MLK-proofs of size polynomial in $n^{s \log_s(n)}$ and depth $2\log_s(n)$.*

Proof: We first outline the idea of the proof. From the antecedent LPHP of PHP_n^{n+1} we immediately derive that for each pigeon i there is at least one variable $p_{i,j}$ that is true. In symbols, $\mathrm{th}_1^{n+1}(p_{i,1}, \ldots, p_{i,n})$. We deduce that among all variables grouped by pigeons, at least $n+1$ are true. In symbols, $\mathrm{th}_{n+1}^{n(n+1)}(p_{1,1}, \ldots, p_{1,n}, \ldots, p_{n+1,1}, \ldots, p_{n+1,n})$. The symmetry of the threshold allows us to show that the same holds when the variables are grouped by holes. In symbols, $\mathrm{th}_{n+1}^{n(n+1)}(p_{1,1}, \ldots, p_{n+1,1}, \ldots, p_{1,n}, \ldots, p_{n+1,n})$. From this, at least one hole contains two pigeons. In symbols, $\mathrm{th}_2^{n+1}(p_{1,i}, \ldots, p_{n+1,i})$ for some $i \in \{1, \ldots, n\}$. This implies RPHP.

According to Lemma 8, we may assume that $n+1$ is an exact power of s since there always is such a number between n and sn. So let us assume that $n = s^r - 1$ for some $r \in \mathbb{N}$. For technical reasons in the proof we will consider a *squared* form (instead of rectangular form) of PHP_n^{n+1} where we assume the existence of an $(n+1)$-st hole in which no pigeon can go. So, we introduce $n+1$ new symbols $p_{1,n+1}, \ldots, p_{n+1,n+1}$ that will stand for the constant 0. For every $i \in \{1, \ldots, n+1\}$, let $p_i = (p_{i,1}, \ldots, p_{i,n+1})$, and let $q_i = (p_{1,i}, \ldots, p_{n+1,i})$. Hence $q_{n+1} = (0, \ldots, 0)$ is the sequence of $n+1$ zeros. Consider the following four sequents.

$$\mathrm{LPHP} \vdash \textstyle\bigwedge_{i=1}^{n+1} \mathrm{th}_1^{n+1}(p_i) \tag{3.1}$$
$$\textstyle\bigwedge_{i=1}^{n+1} \mathrm{th}_1^{n+1}(p_i) \vdash \mathrm{th}_{n+1}^{(n+1)^2}(p_1, \ldots, p_{n+1}) \tag{3.2}$$
$$\mathrm{th}_{n+1}^{(n+1)^2}(p_1, \ldots, p_{n+1}) \vdash \mathrm{th}_{n+1}^{(n+1)^2}(q_1, \ldots, q_{n+1}) \tag{3.3}$$
$$\mathrm{th}_{n+1}^{(n+1)^2}(q_1, \ldots, q_{n+1}) \vdash \mathrm{RPHP} \tag{3.4}$$

In the next lemmas we show how to prove these sequents in MLK. An MLK-proof of $\mathrm{LPHP} \vdash \mathrm{RPHP}$ will follow by three applications of the cut rule. □

Lemma 9 *Sequent (3.1) has MLK-proofs of size polynomial in $n^{s \log_s(n)}$ and depth $2\log_s(n)$.*

Proof: For each $i \in \{1, \ldots, n+1\}$ derive the sequents $\bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{j=1}^{n} p_{i,j} \vee 0$ using right weakening and right $\vee$-introduction. Then, n right $\wedge$-introductions and n left $\wedge$-introductions give $\mathrm{LPHP} \vdash \bigwedge_{i=1}^{n+1} \mathrm{th}_1^{n+1}(p_i)$ by the definition of LPHP and cuts on part (i) of Lemma 6. □

Lemma 10 *Sequent (3.2) has MLK-proofs of size polynomial in $n^{s \log_s(n)}$ and depth $2\log_s(n)$.*

Proof: Recall that $n+1 = s^r$. Let $N = (n+1)^2$. The idea of this proof is to successively pack the conjuncts of the antecedent into a unique threshold formula, following a complete s-ary tree structure of height $\log_s(n+1) = r$. Let $\Sigma = \{0, \ldots, s-1\}$. For every $w \in \Sigma^r$, let $p^w = p_{\overline{w}}$, where $\overline{w}$ is the position of w in the lexicographical order on Σ^r. Thus, $p^{0^r} = p_1$ and $p^{(s-1)^r} = p_{n+1}$. For every $w \in \Sigma^{<r}$, let $p^w = (p^{w0}, \ldots, p^{w(s-1)})$. Observe that $p^\lambda = (p_1, \ldots, p_{n+1})$, where λ is the empty word. For each $t \in \{1, \ldots, r\}$, we exhibit an MLK-proof of

$$\bigwedge_{w \in \Sigma^t} \mathrm{th}^{N/s^t}_{(n+1)/s^t}(p^w) \vdash \bigwedge_{w \in \Sigma^{t-1}} \mathrm{th}^{N/s^{t-1}}_{(n+1)/s^{t-1}}(p^w). \tag{3.5}$$

Once we have all these proofs, we only have to cut sequentially to obtain the lemma. We prove sequent (3.5). For a fixed $t \in \{1, \ldots, r\}$ and a fixed $w \in \Sigma^{t-1}$, an application of part (iii) of Lemma 2 gives

$$\bigwedge_{i=0}^{s-1} \mathrm{th}^{N/s^t}_{(n+1)/s^t}(p^{wi}) \vdash \mathrm{th}^{N/s^{t-1}}_{(n+1)/s^{t-1}}(p^w).$$

We put all these formulas in a unique conjunction using $\wedge$-introduction to get sequent (3.5). □

Lemma 11 *Sequent (3.3) has MLK-proofs of size polynomial in* $n^{s \log_s(n)}$ *and depth* $2\log_s(n)$.

Proof: Immediate from Theorem 2 because $q_1, \ldots, q_{n+1}$ is a permutation of $p_1, \ldots, p_{n+1}$. □

Lemma 12 *Sequent (3.4) has MLK-proofs of size polynomial in* $n^{s \log_s(n)}$ *and depth* $2\log_s(n)$.

Proof: The idea of this proof is to unfold the threshold formula in the antecedent into disjunctions of threshold formulas computing the number of pigeons going into each hole. The unpacking process follows the structure of a complete s-ary tree of height $\log_s(n+1) = r$ in reverse order of that of Lemma 10. We use properties (i) and (ii) of Lemma 7 to perform this process.

Recall that $n + 1 = s^r$. Let $N = (n+1)^2$. Let $\Sigma = \{0, \ldots, s-1\}$. Define $q^w = q_{\overline{w}}$ for every $w \in \Sigma^r$, where $\overline{w}$ is defined as in the proof of Lemma 10. For every $w \in \Sigma^{<r}$, define $q^w = (q^{w0}, \ldots, q^{w(s-1)})$. Observe that $q^\lambda = (q_1, \ldots, q_{n+1})$. For every $t \in \{0, \ldots, r-1\}$ and $w \in \Sigma^t$, properties (ii) and (i) of Lemma 7 give

$$\mathrm{th}^{N/s^t}_{(n+1)/s^t}(q^w) \vdash \mathrm{th}^{N/s^{t+1}}_{(n+1)/s^{t+1}+1}(q^{w0}), \ldots, \mathrm{th}^{N/s^{t+1}}_{(n+1)/s^{t+1}+1}(q^{w(s-2)}), \mathrm{th}^{N/s^{t+1}}_{(n+1)/s^{t+1}}(q^{w(s-1)})$$
$$\mathrm{th}^{N/s^t}_{(n+1)/s^t+1}(q^w) \vdash \mathrm{th}^{N/s^{t+1}}_{(n+1)/s^{t+1}+1}(q^{w0}), \ldots, \mathrm{th}^{N/s^{t+1}}_{(n+1)/s^{t+1}+1}(q^{w(s-1)}).$$

Appropriate cuts and the definition of q^w for $w \in \Sigma^r$ show then that

$$\mathrm{th}^N_{n+1}(q^\lambda) \vdash \mathrm{th}^{n+1}_2(q_1), \mathrm{th}^{n+1}_2(q_2), \ldots, \mathrm{th}^{n+1}_2(q_n), \mathrm{th}^{n+1}_1(q_{n+1}).$$

Since $q_{n+1} = (0, \ldots, 0)$, we immediately have that $\mathrm{th}^{n+1}_1(q_{n+1}) \vdash 0$ by part (i) of Lemma 6, so that the result follows by a cut on $0 \vdash$, successive cuts on part (ii) of Lemma 6, and right $\vee$-introduction. The size of the proof is again quasipolynomial in n. □

Setting $s = 2$ and $s = n^{2/d}$ in Theorem 3, we obtain the main results of this section.

Corollary 1 *The sequent* PHP^{n+1}_n *has MLK-proofs of size* $n^{O(\log n)}$.

Corollary 2 *The sequent* PHP_n^{n+1} *has depth-*d *MLK-proofs of size* $2^{O(n^{3/d})}$ *for every constant* $d > 2$.

Corollary 2 obviously holds for (non-monotone) bounded-depth LK too. The lower bound for the size of depth-d LK proofs of the Pigeonhole Principle is $2^{\Omega(n^{1/6^d})}$ [PBI93, KPW95]. Thus, the dependence on d is an exponential higher than in Corollary 2. This makes a noticeable difference: Corollary 2 implies that there are proofs of quasipolynomial size and depth $O(\log(n)/\log\log(n))$; the lower bound implies only that proofs of quasipolynomial size must have depth $\Omega(\log\log(n))$. It would be interesting to narrow this gap.

3.3 The General Simulation

It is not by chance that the Pigeonhole Principle has quasipolynomial-size MLK-proofs. Indeed, we show that any monotone sequent with LK-proofs of size m has MLK-proofs of size $m^{O(\log(m))}$. We also show that the same result holds in the refutational version of the system.

We say that a formula is in De Morgan normal form if all the negations occur in front of the variables. For every formula φ, let $p(\varphi)$ be a formula in De Morgan normal form that is equivalent to φ. Observe that $p(\varphi)$ is uniformly obtained from φ by pushing the negations to the atoms according to the De Morgan rules. Observe that $p(\neg\neg\varphi) = p(\varphi)$, and that the size of $p(\varphi)$ is linear in the size of φ. We let LK-De Morgan be the subsystem of LK restricted to formulas in De Morgan normal form. In particular, the negation rules are only allowed over variables.

Recall that a proof is called tree-like if every sequent is used at most once as a premise of a rule.

Lemma 13 *The sequents* $\vdash p(\varphi), p(\neg\varphi)$ *and* $p(\varphi), p(\neg\varphi) \vdash$ *have tree-like LK-De Morgan proofs of size quadratic in the size of* φ.

Proof: The proof is by induction on the structure of φ. If φ is atomic, say x, then the sequents $\vdash x, \neg x$ and $x, \neg x \vdash$ are derivable in one step from the axiom $x \vdash x$. Suppose next that φ is of the form $\psi \wedge \theta$. By induction hypothesis, the sequents $\vdash p(\psi), p(\neg\psi)$ and $\vdash p(\theta), p(\neg\theta)$ have tree-like LK-De Morgan proofs of size quadratic in the sizes of ψ and θ respectively. By means of weakening we derive $\vdash p(\psi), p(\neg\psi), p(\neg\theta)$ and $\vdash p(\theta), p(\neg\psi), p(\neg\theta)$. Right $\wedge$-introduction followed by right $\vee$-introduction gives $\vdash p(\psi) \wedge p(\theta), p(\neg\psi) \vee p(\neg\theta)$. The size of the proof is clearly quadratic in the size of φ. The sequent $p(\psi) \wedge p(\theta), p(\neg\psi) \vee p(\neg\theta) \vdash$ is derived similarly. When φ is of the form $\psi \vee \theta$ reason dually. Finally, suppose that φ is of the form $\neg\psi$. By induction hypothesis, the sequent $\vdash p(\psi), p(\neg\psi)$ has a tree-like LK-De Morgan proof of size quadratic in the size of ψ. Since $p(\neg\neg\psi) = p(\psi)$, we immediately have a tree-like LK-De Morgan proof of $\vdash p(\neg\psi), p(\neg\neg\psi)$ of the same size. Reason similarly for the sequent $p(\neg\psi), p(\neg\neg\psi) \vdash$. □

In the following, if Σ is a sequence of formulas $\varphi_1, \ldots, \varphi_r$, we let $p(\Sigma)$ be the sequence $p(\varphi_1), \ldots, p(\varphi_r)$. If $\mathcal{D}$ is a set of sequents, we let $p(\mathcal{D})$ be the result of applying p to all formulas in $\mathcal{D}$.

Theorem 4 *Let Σ and Γ be sequences of formulas and $\mathcal{D}$ a set of sequents. If $\Sigma \vdash \Gamma$ has a tree-like LK-proof from $\mathcal{D}$ of size S, then $p(\Sigma) \vdash p(\Gamma)$ has a tree-like LK-De Morgan proof from $p(\mathcal{D})$ of size $S^{O(1)}$.*

Proof: Suppose that $\Sigma \vdash \Gamma$ has a tree-like LK-proof P from $\mathcal{D}$ of size S. Consider the following transformation of P. First, replace each formula φ in P by $p(\varphi)$. For each right $\neg$-introduction rule in P of the form

$$\frac{\Sigma', \varphi \vdash \Gamma'}{\Sigma' \vdash \neg\varphi, \Gamma',}$$

we simulate the inference

$$\frac{p(\Sigma'), p(\varphi) \vdash p(\Gamma')}{p(\Sigma') \vdash p(\neg\varphi), p(\Gamma')}$$

in the new proof by means of a cut with $\vdash\ p(\varphi), p(\neg\varphi)$, which can be derived in $O(S^2)$ steps according to Lemma 13. Similarly, each left $\neg$-introduction rule in P is replaced by an inference involving a cut with $p(\varphi), p(\neg\varphi) \vdash$. The rest of inferences remain valid as one can easily check. The size of the new proof is clearly $S^{O(1)}$. $\Box$

The following result will be the crucial step of the simulation. Suppose for a moment that in a specified set of truth assignments A, each negated variable $\neg x_i$ is equivalent to a monotone formula φ_i. In such a case, we can treat φ_i as a *pseudocomplement* of x_i with respect to A, and reduce each non-monotone formula to a monotone formula. The interpretation will remain sound as soon as we keep within the set A of truth assignments. Suppose next that $\Sigma \vdash \Gamma$ is a sequent, and that the sequents $\Sigma, \neg x_i \vdash \varphi_i, \Gamma$ and $\Sigma, \varphi_i \vdash \neg x_i, \Gamma$ hold. In such a case, φ_i is the pseudocomplement of x_i with respect to the set of truth assignments that satisfy all formulas of Σ and falsify all formulas of Γ. In this situation, we can turn non-monotone proofs into monotone proofs.

Theorem 5 *Let Σ and Γ be sequences of monotone formulas and $\mathcal{D}$ a set of monotone sequents with all variables within $x_1, \ldots, x_n$. Suppose that for every $i \in \{1, \ldots, n\}$ there exists a monotone formula φ_i such that the sequents $\Sigma \vdash x_i, \varphi_i, \Gamma$ and $\Sigma, \varphi_i, x_i \vdash \Gamma$ have tree-like MLK-proofs from $\mathcal{D}$ of size at most R. Then, if $\Sigma \vdash \Gamma$ has a tree-like LK-proof from $\mathcal{D}$ of size S, then it has a tree-like MLK-proof from $\mathcal{D}$ of size $RS^{O(1)}$.*

Proof: Suppose that $\Sigma \vdash \Gamma$ has a tree-like LK-proof from $\mathcal{D}$ of size S. Since Σ and Γ are sequences of monotone formulas, we have that $p(\Sigma) = \Sigma$ and $p(\Gamma) = \Gamma$. Similarly, $p(\mathcal{D}) = \mathcal{D}$. Therefore, by Theorem 4, the sequent $\Sigma \vdash \Gamma$ has a tree-like LK-De Morgan proof P from $\mathcal{D}$ of size $S^{O(1)}$.

Consider the following transformation on P. First, add Σ to the left of each sequent and Γ to the right of each sequent by weakening on the axioms and initial sequents. Then, replace each occurrence of $\neg x_i$ in P by φ_i. It remains to see how to simulate the rules of $\neg$-introduction. Consider such an application in P

$$\frac{\Sigma', x_i \vdash \Gamma'}{\Sigma' \vdash \neg x_i, \Gamma'.}$$

We need to simulate the inference

$$\frac{\Sigma, \Sigma', x_i \vdash \Gamma', \Gamma}{\Sigma, \Sigma' \vdash \varphi_i, \Gamma', \Gamma.}$$

This is straightforward: derive $\Sigma \vdash x_i, \varphi_i, \Gamma$ and cut on x_i. The simulation of a left $\neg$-introduction rule is symmetrical by means of a cut with $\Sigma, \varphi_i, x_i \vdash \Gamma$. The size of the new proof is clearly $RS^{O(1)}$. $\square$

Recall the following definitions and lemmas from Section 3.1. For every n and $k \in \{0, \ldots, n\}$, let $\mathrm{TH}^n_k : \{0,1\}^n \to \{0,1\}$ be the Boolean function such that $\mathrm{TH}^n_k(a_1, \ldots, a_n) = 1$ if and only if $\sum_{i=1}^k a_i \geq k$, for every $(a_1, \ldots, a_n) \in \{0,1\}^n$. Each TH^n_k is called a threshold function.

Monotone threshold formulas are defined in the following way: $\mathrm{th}^1_0(x) = 1$, $\mathrm{th}^1_1(x) = x$, $\mathrm{th}^1_k(x) = 0$ for every $k > 1$, and for every $n > 1$ and $k \geq 0$, define the formula

$$\mathrm{th}^n_k(x_1, \ldots, x_n) \quad := \bigvee_{(i,j) \in I^n_k} (\mathrm{th}^{n/2}_i(x_1, \ldots, x_{n/2}) \wedge \mathrm{th}^{n-n/2}_j(x_{n/2+1}, \ldots, x_n)),$$

where $I^n_k = \{(i,j) : 0 \leq i \leq n/2,\ 0 \leq j \leq n - n/2,\ i + j \geq k\}$ and $n/2$ is an abbreviation for $\lfloor n/2 \rfloor$. It is straightforward to prove that $\mathrm{th}^n_k(x_1, \ldots, x_n)$ computes the Boolean function TH^n_k. On the other hand, it is easy to prove, by induction on n, that the size of $\mathrm{th}^n_k(x_1, \ldots, x_n)$ is bounded by $n^{O(\log n)}$.

The following lemma is essentially the same as Lemma 3 except for the case $m = n + 1$ and the polynomial bound in the number of lines.

Lemma 14 *For every $n, m, l \in \mathbb{N}$ with $0 < m \leq n + 1$ and $0 \leq l \leq n$, the sequent*

$$\mathrm{th}^n_{m-1}(x_1, \ldots, x_l/0, \ldots, x_n) \vdash \mathrm{th}^n_m(x_1, \ldots, x_l/1, \ldots, x_n)$$

has MLK-proofs with $n^{O(1)}$ lines and size $n^{O(\log n)}$.

Proof: If $0 < m \leq n$, the lemma immediately follows from Lemma 3. The polynomial bound in the number of lines is not explicitely stated there, but it is easy to check. Let us consider the case $m = n + 1$. Our goal is to prove the following sequent:

$$\mathrm{th}^n_n(x_1, \ldots, x_l/0, \ldots, x_n) \vdash 0.$$

We proceed by induction on n. The base case $\mathrm{th}_1^1(0) \vdash 0$ is immediate. Suppose next that $n > 1$. Let us assume without loss of generality that $l \leq n/2$, that is, x_l is in the first block of variables. Recall the definition of $\mathrm{th}_n^n(x_1, \ldots, x_n)$:

$$\bigvee_{(i,j)\in I_n^n} \mathrm{th}_i^{n/2}(x_1, \ldots, x_{n/2}) \wedge \mathrm{th}_j^{n-n/2}(x_{n/2+1}, \ldots, x_n).$$

Since I_n^n is the singleton set $\{(n/2, n-n/2)\}$, the formula $\mathrm{th}_n^n(x_1, \ldots, x_l/0, \ldots, x_n)$ is simply

$$\mathrm{th}_{n/2}^{n/2}(x_1, \ldots, x_l/0, \ldots, x_{n/2}) \wedge \mathrm{th}_{n-n/2}^{n-n/2}(x_{n/2+1}, \ldots, x_n).$$

By induction hypothesis on n, we have $\mathrm{th}_{n/2}^{n/2}(x_1, \ldots, x_l/0, \ldots, x_{n/2}) \vdash 0$. Now left weakening and left $\wedge$-introduction gives

$$\mathrm{th}_n^n(x_1, \ldots, x_l/0, \ldots, x_n) \vdash 0$$

as required. A cut with the axiom $0 \vdash \mathrm{th}_{n+1}^n(x_1, \ldots, x_l/1, \ldots, x_n)$ gives the sequent of the lemma. □

The next lemma easily follows from the definitions of the threshold formulas.

Lemma 15 *For every $n, k \in \mathbb{N}$ with $k > n$, the sequents*

(i) $\mathrm{th}_k^n(x_1, \ldots, x_n) \vdash$, *and*

(ii) $\vdash \mathrm{th}_0^n(x_1, \ldots, x_n)$

have tree-like MLK proofs with $n^{O(1)}$ lines and size $n^{O(\log n)}$.

Proof: Part (ii) of the lemma is simply part (i) in Lemma 2. For part (i), simply observe that $I_k^n = \emptyset$ for every $k > n$, and so $\mathrm{th}_k^n(x_1, \ldots, x_n)$ is 0 by convention (the empty disjunction). Thus, $\mathrm{th}_k^n(x_1, \ldots, x_n) \vdash$ is just an axiom for $k > n$. □

For $k, i \in \mathbb{N}$ with $0 \leq k \leq n$ and $1 \leq i \leq n$, the k-pseudocomplement of x_i is, by definition, the monotone formula $\mathrm{th}_k^n(x_1, \ldots, x_i/0, \ldots, x_n)$. The next Lemma guarantees that the hypothesis of Theorem 5 hold for any of the k-pseudocomplement formulas and any monotone sequent $\Sigma \vdash \Gamma$ with variables within $x_1, \ldots, x_n$ such that Σ contains $\mathrm{th}_k^n(x_1, \ldots, x_n)$ and Γ contains $\mathrm{th}_{k+1}^n(x_1, \ldots, x_n)$.

Lemma 16 *For every $k, i \in \mathbb{N}$ with $0 \leq k \leq n$ and $1 \leq i \leq n$ the sequents*

(i) $\mathrm{th}_k^n(x_1, \ldots, x_n) \vdash \mathrm{th}_{k+1}^n(x_1, \ldots, x_n), \mathrm{th}_k^n(x_1, \ldots, x_i/0, \ldots, x_n), x_i$

(ii) $x_i, \mathrm{th}_k^n(x_1, \ldots, x_i/0, \ldots, x_n), \mathrm{th}_k^n(x_1, \ldots, x_n) \vdash \mathrm{th}_{k+1}^n(x_1, \ldots, x_n)$

have tree-like MLK-proofs with $n^{O(1)}$ lines and size $n^{O(\log n)}$.

Proof: The first sequent follows from right weakening on Lemma 1, part (ii). For the second sequent observe that from Lemma 1, part (i) we have

$$x_i, \mathrm{th}^n_{k+1}(x_1, \dots, x_i/1, \dots, x_n) \vdash \mathrm{th}^n_{k+1}(x_1, \dots, x_n).$$

Moreover, $\mathrm{th}^n_k(x_1, \dots, x_i/0, \dots, x_n) \vdash \mathrm{th}^n_{k+1}(x_1, \dots, x_i/1, \dots, x_n)$ by Lemma 14. The sequent in (ii) is obtained by cutting and then adding $\mathrm{th}^n_k(x_1, \dots, x_n)$ by left weakening. □

Theorem 6 *Let $\Sigma \vdash \Gamma$ be a monotone sequent and let $\mathcal{D}$ be a set of monotone sequents with n variables. If $\Sigma \vdash \Gamma$ has an LK-proof from $\mathcal{D}$ of size S, then $\Sigma \vdash \Gamma$ has a tree-like MLK-proof from $\mathcal{D}$ with $S^{O(1)}$ lines and size $S^{O(1)} \cdot n^{O(\log n)}$.*

Proof: By Theorem 4 and the well known result that tree-like LK polynomially simulates LK [Kra95], it will be sufficient to simulate tree-like LK-De Morgan proofs by tree-like MLK-proofs. Let P be a tree-like LK-De Morgan proof of $\Sigma \vdash \Gamma$ from $\mathcal{D}$ of size S. By trivial weakening, the sequent

$$\mathrm{th}^n_k(x_1, \dots, x_n), \Sigma \vdash \Gamma, \mathrm{th}^n_{k+1}(x_1, \dots, x_n)$$

has tree-like LK-De Morgan proofs from $\mathcal{D}$ of size S. Moreover, by weakening on the previous lemma, the formula $\mathrm{th}^n_k(x_1, \dots, x_i/0, \dots, x_n)$ is the pseudocomplement of x_i on sequents having $\mathrm{th}^n_k(x_1, \dots, x_n)$ on the left and $\mathrm{th}^n_{k+1}(x_1, \dots, x_n)$ on the right. It follows by Theorem 5 that for each $k \in \{0, \dots, n\}$ there are tree-like MLK-proofs from $\mathcal{D}$ of the sequents $\mathrm{th}^n_k(x_1, \dots, x_n), \Sigma \vdash \Gamma, \mathrm{th}^n_{k+1}(x_1, \dots, x_n)$ each one with $S^{O(1)}$ lines and size $S^{O(1)} \cdot n^{O(\log n)}$. Finally, n consecutive cuts give us a proof of the sequent $\mathrm{th}^n_0(x_1, \dots, x_n), \Sigma \vdash \Gamma, \mathrm{th}^n_{n+1}(x_1, \dots, x_n)$ from which we obtain the theorem using Lemma 15. □

We say that a proof system S_1 quasipolynomially simulates a proof system S_2 if every S_2-proof of size m can be turned into an S_1-proof of size $m^{O(\log m)}$. We have the following interesting corollary.

Corollary 3 *Tree-like MLK quasipolynomially simulates LK on monotone sequents. In particular, tree-like MLK quasipolynomially simulates MLK.*

The careful reader will notice that the proof of Theorem 6 shows that the number of lines of the resulting MLK proof is polynomial in n and the number of lines of the original LK proof. This observation reveals that any proof of a superpolynomial gap between LK and MLK, if any, should focus on size and not on the number of lines.

Since every MLK-proof can be polynomially simulated by a proof in the intuitionistic calculus JK (see [B01]) we get the following.

Corollary 4 *On monotone sequents, JK quasipolynomially simulates LK.*

Note, however, that this is unlikely for intuitionistically valid *nonmonotone* sequents, see [BP01].

We close this section with a result concerning the refutational version of MLK. A weakness of the simulation result, as stated in Corollary 3, is that it only makes sense on monotone sequents. Nonetheless, the simulation is somewhat stronger as it works for the refutational version of MLK. Hence, since each tautology can be turned into a contradictory set of clauses of size polynomial in the size of the tautology, and each clause can be written as a monotone sequent itself, the result is in a sense general enough to deal with *all* tautologies. To state our result, we have to define the translation of an arbitrary tautology into a contradictory set of clauses precisely.

Let $\Sigma \vdash \Gamma$ be a valid sequent. Here, Σ and Γ are sequences of Boolean formulas $\varphi_1, \ldots, \varphi_r$ and $\psi_1, \ldots, \psi_s$. What this means is that the formula

$$\bigwedge_{i=1}^{r} \varphi_i \wedge \bigwedge_{i=1}^{s} \neg\psi_i$$

is a contradiction. This formula is sometimes written as $\bigwedge \Sigma \wedge \bigwedge \neg\Gamma$. We show how to turn each contradictory Boolean formula into a contradictory set of clauses of polynomial-size in the size of the formula. The translation is quite standard but we define it precisely since it will be needed in a later section.

Given a Boolean formula φ over the basis $\{\wedge, \vee, \neg\}$ with variables $x_1, \ldots, x_n$ and constants 0 and 1, let y_θ be a new propositional variable for each subformula θ of φ (including φ itself). We define a set of clauses $C(\varphi)$ by induction on the construction of φ:

(i) If φ is the constant 0, let $C(\varphi) = \{\neg y_\varphi\}$.

(ii) If φ is the constant 1, let $C(\varphi) = \{y_\varphi\}$.

(iii) If φ is a variable x_i, let $C(\varphi) = \{\neg x_i \vee y_\varphi, x_i \vee \neg y_\varphi\}$.

(iv) If $\varphi = \neg\psi$, let $C(\varphi) = \{y_\psi \vee y_\varphi, \neg y_\psi \vee \neg y_\varphi\} \cup C(\psi)$.

(v) If $\varphi = \psi \wedge \xi$, let $C(\varphi) = \{\neg y_\psi \vee \neg y_\xi \vee y_\varphi, y_\psi \vee \neg y_\varphi, y_\xi \vee \neg y_\varphi\} \cup C(\psi) \cup C(\xi)$.

(vi) If $\varphi = \psi \vee \xi$, let $C(\varphi) = \{\neg y_\psi \vee y_\varphi, \neg y_\xi \vee y_\varphi, y_\psi \vee y_\xi \vee \neg y_\varphi\} \cup C(\psi) \cup C(\xi)$.

The following lemma expresses the main property of $C(\varphi)$.

Lemma 17 *The formula φ is satisfiable if and only if the set of clauses $C(\varphi) \cup \{y_\varphi\}$ is.*

Proof: View φ as a Boolean circuit, and interpret each variable y_θ as giving a truth value to the internal gate that corresponds to the root of θ. Each y_θ is determined by the $x_1, \ldots, x_n$. Moreover, y_θ is the value of the subcircuit θ on input $x_1, \ldots, x_n$. Hence, if σ satisfies φ, then σ can be extended to an assignment σ' that satisfies $C(\varphi) \cup \{y_\varphi\}$. Similarly, if σ' satisfies $C(\varphi) \cup \{y_\varphi\}$, the restriction of σ' to $x_1, \ldots, x_n$ satisfies φ. □

Observe that a clause $x_1 \vee \ldots \vee x_r \vee \neg y_1 \vee \ldots \vee \neg y_s$ can be written as a monotone sequent:

$$y_1, \ldots, y_s \vdash x_1, \ldots, x_r.$$

Now we can state our result in terms of refutations of sets of clauses. We just need to take $\Sigma \vdash \Gamma$ to be the empty sequent in Theorem 6.

Corollary 5 *As refutation systems, tree-like MLK quasipolynomially simulates LK on contradictory sets of clauses.*

According to the observations preceeding the statement of this corollary, this result is more general than Corollary 3.

3.4 Proving a Gap is Hard

In this section we prove the surprising result that, as refutation systems, MLK is polynomially bounded if and only if LK is polynomially bounded. This shows that proving a superpolynomial gap between MLK and LK is at least as hard as proving that LK and all Frege systems are not polynomially bounded. This is a long-standing open problem.

Before we state our result, we need the following lemma due to Pudlák and appearing in [AGP01] with only a sketch of the proof. For completeness, we provide a full proof.

Lemma 18 *Let $\tau_k^n(x_1, \ldots, x_n)$ be polynomial-size monotone formulas for TH_k^n. If the sequents of Lemmas 15 and 16 have polynomial-size LK-proofs with $\tau_k^n(x_1, \ldots, x_n)$ instead of $\mathrm{th}_k^n(x_1, \ldots, x_n)$, then MLK polynomially simulates LK on monotone sequents, even as refutation systems.*

Proof: Let $p(n)$ be a bound on the size of $\tau_k^n(x_1, \ldots, x_n)$. Suppose we have LK-proofs of all the sequents in Lemmas 15 and 16 in a single proof of size at most $q(n)$. These $2(n+1)n+2$ sequents are called the *pseudocomplement properties*. We prove, by induction on n, that all $2(n+1)n+2$ pseudocomplement properties of $\tau_k^n(x_1, \ldots, x_n)$ can be obtained in a single MLK-proof of polynomial-size. This will be enough since then we can apply the same argument as in Theorem 6 with $\tau_k^n(x_1, \ldots, x_n)$ instead of $\mathrm{th}_k^n(x_1, \ldots, x_n)$.

We will obtain a recurrence $s(n)$ for the size of the MLK-proofs of the pseudocomplement properties of τ_k^n. For $n = 1$, the proofs are just constant size and $s(n) = O(1)$. Suppose $n > 1$ next. Define auxiliary formulas as follows. Let $\sigma_0^n = 1$, $\sigma_n^n = \bigwedge_{i=1}^n x_i$, and for every $i \in \{1, \ldots, n-1\}$, let σ_i^n be the formula

$$\tau_i^{n-1}(x_1, \ldots, x_{n-1}) \vee (\tau_{i-1}^{n-1}(x_1, \ldots, x_{n-1}) \wedge x_n).$$

Observe that the size of σ_k^n is bounded by $2p(n-1)+3$. It is easy to get MLK-proofs of the sequents of the pseudocomplement properties for σ_i^n from those of τ_i^{n-1}. We sketch how in Lemma 19 below. The size of these proofs is at most $l(n)$ for some polynomial $l(n)$. With these, we use the argument of Theorem 6 with $\sigma_0^n, \ldots, \sigma_n^n$ as threshold formulas to turn the LK-proofs of the properties of τ_i^n into MLK-proofs of size $q(n) + 2q(n)(2p(n-1)+3) + s(n-1) + l(n)$. To see this bound on the size, observe that the proof is built as follows. For each $k \in \{0, \ldots, n\}$, take the LK-De Morgan proofs of the $2(n+1)n+2$ properties for τ_i^n and add σ_k^n to the left and σ_{k+1}^n to the right by weakening. This gives size $q(n) + 2q(n)(2p(n-1)+3)$. Then replace each negated variable $\neg x_i$ by $\sigma_k^n(x_i/0)$. This gives $q(n)(2p(n-1)+3)$ additional symbols. Then derive the pseudocomplement properties of τ_i^{n-1} monotonically in size $s(n-1)$, and those of σ_k^n from these in $l(n)$ additional symbols. Finally, the rules of $\neg$-introduction are simulated by cuts on these sequents. This analysis gives us the recurrence $s(n) = q(n)+3q(n)(2p(n-1)+3)+s(n-1)+l(n)$ which is easily seen to give a polynomial. We note that the proofs are not tree-like at all. □

For the sake of completeness, we sketch how to get some pseudocomplement properties of σ_k^n from those of τ_k^{n-1}.

Lemma 19 *There exists a polynomial $l(n)$ such that the pseudocomplement property*

$$\sigma_k^n \vdash \sigma_{k+1}^n, \sigma_k^n(x_i/0), x_i$$

has an MLK-proof of size $l(n)$ from the pseudocomplement properties of $\tau_0^{n-1}, \ldots, \tau_{n-1}^{n-1}$. Here, $i \neq n$.

Proof: By weakening on axioms, we trivially have

$$\tau_k^{n-1} \quad \vdash \quad \tau_k^{n-1}, \tau_{k-1}^{n-1} \wedge x_n \tag{3.6}$$

$$\tau_{k-1}^{n-1} \wedge x_n \quad \vdash \quad \tau_k^{n-1}, \tau_{k-1}^{n-1} \wedge x_n. \tag{3.7}$$

By left $\vee$-introduction we get

$$\sigma_k^n \vdash \tau_k^{n-1}, \tau_{k-1}^{n-1} \wedge x_n. \tag{3.8}$$

We have $\tau_k^{n-1} \vdash \tau_{k+1}^{n-1}, \tau_k^{n-1}(x_i/0), x_i$ as an initial sequent. A cut with equation (3.8) gives

$$\sigma_k^n \vdash \tau_{k+1}^{n-1}, \tau_k^{n-1}(x_i/0), x_i, \tau_{k-1}^{n-1} \wedge x_n. \tag{3.9}$$

On the other hand, by weakening on an axiom and right $\vee$-introduction we have $\tau_{k+1}^{n-1} \vdash \sigma_{k+1}^n$. A cut with equation (3.9) gives

$$\sigma_k^n \vdash \sigma_{k+1}^n, \tau_k^{n-1}(x_i/0), x_i, \tau_{k-1}^{n-1} \wedge x_n. \tag{3.10}$$

Again for trivial reasons, we have $\tau_k^{n-1}(x_i/0) \vdash \sigma_k^n(x_i/0)$. A cut with equation (3.10) gives

$$\sigma_k^n \vdash \sigma_{k+1}^n, \sigma_k^n(x_i/0), x_i, \tau_{k-1}^{n-1} \wedge x_n. \tag{3.11}$$

Trivially again, $\tau_{k-1}^{n-1} \wedge x_n \vdash \tau_{k-1}^{n-1}$. Moreover, $\tau_{k-1}^{n-1} \vdash \tau_k^{n-1}, \tau_{k-1}^{n-1}(x_i/0), x_i$ is an initial sequent. Together, a cut gives

$$\tau_{k-1}^{n-1} \wedge x_n \vdash \tau_k^{n-1}, \tau_{k-1}^{n-1}(x_i/0), x_i. \tag{3.12}$$

Still trivially $\tau_k^{n-1} \vdash \sigma_k^n$. Moreover, by Lemma 1 we have $\sigma_k^n \vdash \sigma_k^n(x_i/0), x_i$. A cut gives $\tau_k^{n-1} \vdash \sigma_k^n(x_i/0), x_i$, and a cut with equation (3.12) gives

$$\tau_{k-1}^{n-1} \wedge x_n \vdash \sigma_k^n(x_i/0), x_i, \tau_{k-1}^{n-1}(x_i/0), x_i. \tag{3.13}$$

Trivially again, $\tau_{k-1}^{n-1} \wedge x_n \vdash x_n$. Right introduction of conjunction and contraction gives then

$$\tau_{k-1}^{n-1} \wedge x_n \vdash \sigma_k^n(x_i/0), x_i, \tau_{k-1}^{n-1}(x_i/0) \wedge x_n. \tag{3.14}$$

Again $\tau_{k-1}^{n-1}(x_i/0) \wedge x_n \vdash \sigma_k^n(x_i/0)$ trivially. A cut with equation (3.14) and contraction gives $\tau_{k-1}^{n-1} \wedge x_n \vdash \sigma_k^n(x_i/0), x_i$. Finally, a cut with equation (3.11) and contraction gives

$$\sigma_k^n \vdash \sigma_{k+1}^n, \sigma_k^n(x_i/0), x_i. \tag{3.15}$$

This is exactly what we needed. □

Back to the problem of this section, the lemma we just proved will help us prove the first half of our result. For the other half of the proof, we need another technical lemma. Intuitively, this lemma will say that it is enough to have short LK-refutations of sets of clauses in order to have short LK-proofs of arbitrary sequents.

Recall the definition of the set of clauses $C(\varphi)$ for every Boolean formula φ in Section 3.3. Recall also the notation $\bigwedge \Sigma \wedge \bigwedge \neg\Gamma$ defined there.

Lemma 20 *Let $\Sigma \vdash \Gamma$ be a sequent, let φ be the formula $\bigwedge \Sigma \wedge \bigwedge \neg\Gamma$, and let $\mathcal{D}$ be the set of sequents corresponding to the set of clauses $C(\varphi) \cup \{y_\varphi\}$. If $\mathcal{D}$ has an LK-refutation of size S, then $\Sigma \vdash \Gamma$ has an LK-proof of size $S^{O(1)}$.*

Proof: Let P be an LK-refutation of $\mathcal{D}$. Replace each variable y_θ, with θ a subformula of φ, by the formula θ itself. The result is still a valid refutation from a set of sequents. Moreover, we claim that all clauses in $C(\varphi)$ become trivially derivable (note that we do not claim this for the additional clause y_φ). Indeed, if $C \in C(\varphi)$ is a clause that comes from a constant, this is trivial since it is the clause 1. If C is a clause that comes from a variable x_i, the resulting sequent is $x_i \vdash x_i$, again trivially derivable. If C comes from $\neg\psi$, the resulting sequent is $\psi \vdash \psi$, again trivial. If C comes from $\psi \wedge \xi$, the resulting sequent is one of $\psi, \xi \vdash \psi \wedge \xi$, $\psi \wedge \xi \vdash \psi$ or $\psi \wedge \xi \vdash \xi$. All three cases are trivial to derive. Similarly if C comes from $\psi \vee \xi$. Thus, what we really have is a refutation of the sequent $\vdash \varphi$. Now build the proof of $\Sigma \vdash \Gamma$ as follows.

Assume for simplicity that Σ and Γ are single formulas σ and γ. The general case is similar. Start with the axioms $\sigma \vdash \sigma$ and $\neg\gamma \vdash \neg\gamma$. Apply right $\wedge$-introduction to obtain $\sigma, \neg\gamma \vdash \sigma \wedge \neg\gamma$. Observe that $\sigma \wedge \neg\gamma$ is φ. Simulate the LK-refutation of $\vdash \varphi$ on $\sigma, \neg\gamma \vdash \varphi$. This gives a proof of the sequent $\sigma, \neg\gamma \vdash$. Finally, a cut with $\vdash \gamma, \neg\gamma$ gives the sequent $\sigma \vdash \gamma$. The size of this LK-proof is obviously $S^{O(1)}$. $\square$

The following is a converse to the previous Lemma.

Lemma 21 *Let $\mathcal{D} = \{\Sigma_i \vdash \Gamma_i : i = 1, \ldots, r\}$ be a set of sequents and for every $i \in \{1, \ldots, r\}$ let φ_i be the formula $\bigwedge \Sigma_i \wedge \bigwedge \neg\Gamma_i$. If the sequent $\vdash \bigvee_i \varphi_i$ has an LK-proof of size S, then $\mathcal{D}$ has an LK-refutation of size $S^{O(1)}$.*

Proof: Start with each initial sequent $\Sigma_i \vdash \Gamma_i$ and derive the sequent $\bigwedge \Sigma_i \wedge \bigwedge \neg\Gamma_i \vdash$ by left introduction of $\neg$, and left introduction of $\wedge$. Then apply left introduction of $\vee$ to obtain the sequent $\bigvee_i \varphi_i \vdash$. Finally, derive the sequent $\vdash \bigvee_i \varphi_i$ in size S and cut to obtain the empty sequent. $\square$

We are now ready to prove the main result of this section.

Theorem 7 *As refutation systems, LK is polynomially bounded if and only if MLK is polynomially bounded.*

Proof: Suppose that LK is polynomially bounded as a refutation system. By Lemma 20, it is also polynomially bounded as a *direct* proof system. Let τ_k^n be Valiant's monotone formulas for all threshold functions [Val84]. Any other polynomial-size monotone formulas computing TH_k^n would do as well. Since we are assuming that LK is polynomially bounded, the pseudocomplement properties of τ_k^n have polynomial-size LK-proofs. Hence, by Lemma 18, MLK polynomially simulates LK even as a refutation system. In particular, MLK is polynomially bounded as a refutation system.

For the other direction, suppose that MLK is polynomially bounded as a refutation system. Let $\mathcal{D}$ be a contradictory set of sequents and let $\varphi = \bigvee_i \varphi_i$ be the formula as in Lemma 21. Note that $\vdash \varphi$ is a valid sequent since φ is a tautology. We obtain the set of clauses $C(\neg\varphi) \cup \{y_\varphi\}$. Let $\mathcal{D}'$ be the resulting set of monotone sequents. Since MLK is polynomially bounded as a refutation system, $\mathcal{D}'$ has a polynomial-size MLK-refutation. It follows by Lemma 20 that $\vdash \varphi$ has a polynomial-size LK-proof. Finally, apply Lemma 21 to obtain an LK-refutation of $\mathcal{D}$ of polynomial-size. $\square$

3.5 Some Specific Tautologies and Separation Results

The purpose of this section is to prove that certain monotone sequents of interest, such as the Functional Pigeonhole Principle, the Onto Pigeonhole Principle and the Matching Principle admit polynomial-size monotone proofs. In order to prove these upper bounds, we realize that each of the variables has a polynomial-size monotone formula for its pseudocomplement function. This will allow us simulate their polynomial-size Frege proofs with only a polynomial overhead in size. We also consider the Clique-Coclique Principle, and in this case, we do a monotone reduction to the Functional Pigeonhole Principle in order to obtain polynomial-size monotone proofs.

As corollaries to these results, we obtain exponential separations between MLK and Cutting Planes, and MLK and Bounded-depth LK.

3.5.1 Pigeonhole Principles

As we already know, the PHP expresses the fact that there cannot be a one-to-one *correspondence* from a set of $n+1$ pigeons into a set of n holes. The term *correspondence* is used to emphasise that we do not insist that pigeons be mapped to a unique hole. The Functional PHP, instead, expresses the fact that there cannot be a one-to-one *function* from a set of $n+1$ pigeons into a set of n holes. Thus, the Functional PHP is a weaker statement. Finally, the Onto PHP expresses the fact that there cannot be a one-to-one correspondence from a set of $n+1$ pigeons *onto* a set of n holes. We could also consider the Functional and Onto PHP expressing that there cannot be a one-to-one functions from $n+1$ pigeons onto n holes.

We consider the following propositional formulations of PHP, Onto PHP and Functional PHP: We let PHP_n^{n+1} be the sequent

$$\bigwedge_{i=1}^{n+1} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i\neq j}}^{n+1} (p_{i,k} \wedge p_{j,k}).$$

We let OPHP_n^{n+1} be the sequent

$$\bigwedge_{i=1}^{n+1} \bigvee_{j=1}^{n} p_{i,j} \wedge \bigwedge_{j=1}^{n} \bigvee_{i=1}^{n+1} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i\neq j}}^{n+1} (p_{i,k} \wedge p_{j,k}).$$

Finally, we let FPHP_n^{n+1} be the sequent

$$\bigwedge_{i=1}^{n+1} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i\neq j}}^{n+1} (p_{i,k} \wedge p_{j,k}) \vee \bigvee_{k=1}^{n+1} \bigvee_{\substack{i,j=1 \\ i\neq j}}^{n} (p_{k,i} \wedge p_{k,j}).$$

Using Corollary 4 and Buss' polynomial size LK proofs of the PHP we give another proof of the main result of Section 3.2.

Theorem 8 PHP_n^{n+1} *has MLK-proofs of size quasipolynomial in* n.

We can improve this result showing that the principles OPHP, FPHP and a Perfect Matching Principle PM that we introduce later admit polynomial size MLK proofs.

Theorem 9 FPHP_n^{n+1} *and* OPHP_n^{n+1} *have tree-like MLK-proofs of size polynomial in* n.

Proof: Buss proved that PHP_n^{n+1} has a Frege proof of size polynomial in n, and therefore, so do FPHP_n^{n+1} and OPHP_n^{n+1}. Since tree-like LK polynomially simulates any Frege system [Kra95], they also have polynomial-size tree-like LK-proofs. We first consider FPHP_n^{n+1}. For every $i \in \{1, \ldots, n+1\}$ and $j \in \{1, \ldots, n\}$, let φ_{ij} be the formula $\bigvee_{j' \neq j} p_{i,j'}$ where j' ranges over $\{1, \ldots, n\}$. Let LFPHP_n^{n+1} be the left hand side of the sequent FPHP_n^{n+1}, and let RFPHP_n^{n+1} be the right hand side of the sequent FPHP_n^{n+1}. We claim that the sequents

$$\mathrm{LFPHP}_n^{n+1} \vdash p_{i,j}, \varphi_{ij}, \mathrm{RFPHP}_n^{n+1} \tag{3.16}$$

$$\mathrm{LFPHP}_n^{n+1}, \varphi_{ij}, p_{i,j} \vdash \mathrm{RFPHP}_n^{n+1} \tag{3.17}$$

have tree-like MLK-proofs of size polynomial in n. The result will follow for FPHP_n^{n+1} by Theorem 5. For sequent (3.16) reason as follows. For every $j' \in \{1, \ldots, n\}$, we have $p_{i,j'} \vdash p_{i,1}, \ldots, p_{i,n}, \mathrm{RFPHP}_n^{n+1}$ by right weakening on the axiom $p_{i,j'} \vdash p_{i,j'}$ and structural rules. By left $\vee$-introduction we get $\bigvee_{j=1}^n p_{i,j} \vdash p_{i,1}, \ldots, p_{i,n}, \mathrm{RFPHP}_n^{n+1}$. Left weakening and left $\wedge$-introduction gives $\mathrm{LFPHP}_n^{n+1} \vdash p_{i,1}, \ldots, p_{i,n}, \mathrm{RFPHP}_n^{n+1}$. Finally, some structural rules and right $\vee$-introduction give sequent (3.16). For sequent (3.17) reason as follows. For every $j, j' \in \{1, \ldots, n+1\}$ such that $j \neq j'$, we have $p_{i,j}, p_{i,j'} \vdash p_{i,j} \wedge p_{i,j'}$ easily. Left weakening, right weakening and right $\vee$-introduction gives $\mathrm{LFPHP}_n^{n+1}, p_{i,j}, p_{i,j'} \vdash \mathrm{RFPHP}_n^{n+1}$. Finally, left $\vee$-introduction for every $j' \neq j$ gives sequent (3.17). As regards OPHP_n^{n+1}, one simply needs define φ_{ij} as $\bigvee_{i' \neq i} p_{i',j}$ where i' ranges over $\{1, \ldots, n+1\}$, and reason analogously. □

3.5.2 Matching Principle

Let us be given a graph $G = (V, E)$ on $n = 3m$ nodes. We consider the following matching principle PM_n formulated in [IPU94]. If X is a set of m edges forming a perfect matching in G and Y is an $m-1$ subset of V, then there is some edge $(u, v) \in X$ such that neither u nor v are in V. To encode this principle as a monotone sequent we use variables x_{ik} for $i \in \{1, \ldots, m\}$ and $k \in \{1, \ldots, 3m\}$ whose intended meaning is that the node k is in the i-th edge of the matching, and variables y_{ik} for $i \in \{1, \ldots, m-1\}$ and $k \in \{1, \ldots, 3m\}$ whose intended meaning is that the node k is NOT the i-th element in Y. We will encode the fact that there is a perfect matching on

m edges in G by an $m \times 3m$ matrix such that in each row there are exactly two 1's and in each column there is at most one 1. Notice that our formula has depth 3.

$$\begin{aligned}
X_1 &\equiv \bigwedge_{i=1}^{m} \bigvee_{\substack{k,k'=1\\k\neq k'}}^{3m} (x_{ik} \wedge x_{ik'}) \\
X_2 &\equiv \bigwedge_{i=1}^{m} \bigwedge_{\substack{k,l,h=1\\k\neq l\neq h\neq k}}^{3m} (\neg x_{ik} \vee \neg x_{il} \vee \neg x_{ih}) \\
X_3 &\equiv \bigwedge_{\substack{i,i'=1\\i\neq i'}}^{m} \bigwedge_{k=1}^{3m} (\neg x_{ik} \vee \neg x_{i'k}).
\end{aligned}$$

Similarly, we will encode that Y is an $m-1$ subset of V, by an $(m-1) \times 3m$ matrix in which for each row there is exactly one 0 and in each column there is at most one 0 (recall that the presence of a node in Y is indicated by a negated variable).

$$\begin{aligned}
Y_1 &\equiv \bigwedge_{\substack{i,i'=1\\i\neq i'}}^{m-1} \bigwedge_{k=1}^{3m} (y_{ik} \vee y_{i'k}) \\
Y_2 &\equiv \bigwedge_{i=1}^{m-1} \bigwedge_{\substack{k,k'=1\\k\neq k'}}^{3m} (y_{ik} \vee y_{ik'}) \\
Y_3 &\equiv \bigwedge_{i=1}^{m-1} \bigvee_{k=1}^{3m} \neg y_{ik}
\end{aligned}$$

The last formula that we introduce means that there is an edge such that neither of its endpoints is in Y.

$$XY \equiv \bigvee_{i=1}^{m} \bigvee_{\substack{k,k'=1\\k\neq k'}}^{3m} \left(x_{ik} \wedge x_{ik'} \wedge \bigwedge_{i=1}^{m-1} y_{ik} \wedge \bigwedge_{i=1}^{m-1} y_{ik'} \right)$$

Then, the PM_{3m} principle is expressed by the following sequent:

$$X_1, X_2, X_3, Y_1, Y_2, Y_3 \vdash XY \tag{3.18}$$

We turn this into a monotone sequent. Consider the formulas $X_i^D \equiv \neg X_i$ for $i = 2, 3$ and $Y_3^D \equiv \neg Y(3)$, where $\neg X_i$ and $\neg Y_i$ stands for the result of switching $\wedge$ and $\vee$, and replacing $\neg v$ by v for every variable. Then, the sequent 3.18 is equivalent to the monotone sequent

$$X_1, Y_1, Y_2 \vdash X_2^D, X_3^D, Y_3^D, XY$$

Notice that, as observed in [IPU94], PM_{3m} can be reduced to OPHP^m_{m-1}. However we need to define the PHP variables p_{ij} as $\bigvee_{k=1}^{3m}(x_{ik} \wedge \neg y_{jk})$ which is not a monotone formula. Therefore the reduction cannot be proved in MLK. Either way we can get polynomial size MLK-proofs for the PM_{3m} principle directly.

Theorem 10 PM_n *has tree-like MLK-proofs of size polynomial in* n.

Proof: Since [IPU94] gave polynomial-size LK proofs of PM_n, it will suffice to define pseudocomplement formulas for x_{ik} and y_{jk}. Define for each $i \in \{1, \ldots, m\}$ and for each $k \in \{1, \ldots, 3m\}$ the pseudocomplement formula $\varphi(x_{ik})$ for x_{ik} as:

$$\bigvee_{\substack{k', k''=1 \\ k \neq k' \neq k'' \neq k}}^{3m} (x_{ik'} \wedge x_{ik''}).$$

For each $i \in \{1, \ldots, m-1\}$ and each $k \in \{1, \ldots, 3m\}$ define the pseudocomplement formula $\varphi(y_{ik})$ for y_{ik} as

$$\bigwedge_{\substack{k'=1 \\ k' \neq k}}^{3m} y_{ik'}.$$

We prove that for each $i \in \{1, \ldots, m\}$, $j \in \{1, \ldots, m-1\}$ and $k \in \{1, \ldots, 3m\}$ the following sequents have polynomial size tree-like MLK-proofs:

$$X_1, x_{ik}, \varphi(x_{ik}) \vdash X_2^D, X_3^D \tag{3.19}$$

$$X_1 \vdash x_{ik}, \varphi(x_{ik}), X_2^D, X_3^D \tag{3.20}$$

$$Y_1, Y_2, y_{jk}, \varphi(y_{jk}) \vdash Y_3^D \tag{3.21}$$

$$Y_1, Y_2 \vdash y_{jk}, \varphi(y_{jk}), Y_3^D \tag{3.22}$$

We prove sequents (3.19) and (3.20). Sequents (3.21) and (3.22) follow by an argument similar to that of FPHP. Observe that X_2^D, X_3^D are the following formulas

$$\bigvee_{i=1}^{m} \bigvee_{\substack{k,l,h=1 \\ k \neq l \neq h \neq k}}^{3m} (x_{ik} \wedge x_{il} \wedge x_{ih})$$

and

$$\bigvee_{\substack{i,i'=1 \\ i \neq i'}}^{m} \bigvee_{k=1}^{3m} (x_{ik} \wedge x_{i'k})$$

For sequent (3.19) reason as follows: for each $k' \neq k'' \in \{1, \ldots, 3m\}$, $k', k'' \neq k$, we have proofs of the sequents $x_{ik}, (x_{ik'} \wedge x_{ik''}) \vdash (x_{ik} \wedge x_{ik'} \wedge x_{ik''})$. By weakenings and right $\vee$-introduction we obtain $x_{ik}, (x_{ik'} \wedge x_{ik''}) \vdash X_3^D$. By right $\vee$-introductions on all previous proofs we have $x_{ik}, \varphi(x_{ik}) \vdash X_3^D$ from which sequent (3.19) follows by two weakenings, left and right.

For sequent (3.20) reason as follows: for each $k' \neq k$ we have proofs of the sequents $x_{ik} \wedge x_{ik'} \vdash x_{ik}$. By left $\vee$-introduction we can derive

$$\bigvee_{\substack{k'=1 \\ k' \neq k}}^{3m} (x_{ik} \wedge x_{ik'}) \vdash x_{ik}.$$

From this, by right weakening we have

$$\bigvee_{\substack{k'=1\\k'\neq k}}^{3m}(x_{ik}\wedge x_{ik'})\vdash x_{ik},\varphi(x_{ik}) \tag{3.23}$$

For each $k',k''\in\{1,\ldots,3m\}$, with $k\neq k'\neq k''\neq k$ we can derive $x_{ik'}\wedge x_{ik''}\vdash x_{ik'}\wedge x_{ik''}$. From this, by right weakenings, we can derive $x_{ik'}\wedge x_{ik''}\vdash x_{ik},\varphi(x_{ik})$. By left $\vee$-introductions on these proofs we obtain

$$\bigvee_{\substack{k',k''=1\\k\neq k'\neq k''\neq k}}^{3m}(x_{ik'}\wedge x_{ik''})\vdash x_{ik},\varphi(x_{ik}) \tag{3.24}$$

Finally by left $\vee$-introduction between (3.23) and (3.24), left weakening, and left $\wedge$-introduction we obtain $X_1\vdash x_{ik},\varphi(x_{ik})$, from which (3.20) follows by right weakenings. □

3.5.3 Clique-Coclique Principle

A graph G is a k-clique if there is a set of k nodes of G such that any two distinct nodes of the set are connected by an edge, and no other edge is present in G. A graph G is a k-coclique if there is a partition of the nodes of G into k disjoint sets in such a way that any two nodes that belong to different sets are connected by an edge, and no other edges are present in G.

The (n,k)-clique-coclique principle of [BPR97] says that, given a set V of n nodes, if G is a k-clique over V and H is a $(k-1)$-coclique over V, then there is an edge in G that is not present in H. This principle may be stated as a monotone sequent CLIQUE_k^n as follows. For every $l\in\{1,\ldots,k\}$ and $i\in\{1,\ldots,n\}$, let x_{li} be a propositional variable whose intended meaning is that i is the l-th node in the fully connected set of a k-clique over $\{1,\ldots,n\}$. Similarly, for every $t\in\{1,\ldots,k-1\}$ and $i\in\{1,\ldots,n\}$, let y_{it} be a propositional variable whose intended meaning is that the i-th node is in the t-th disjoint set of a $(k-1)$-coclique over $\{1,\ldots,n\}$. The principle is then expressed as follows

$$\bigwedge_{l=1}^{k}\bigvee_{i=1}^{n}x_{li}\wedge\bigwedge_{i=1}^{n}\bigvee_{t=1}^{k-1}y_{it}\vdash\bigvee_{t=1}^{k-1}\bigvee_{\substack{l,l'=1\\l\neq l'}}^{k}\bigvee_{\substack{i,j=1\\i\neq j}}^{n}(x_{li}\wedge x_{l'j}\wedge y_{it}\wedge y_{jt})\vee\mathrm{BAD}$$

where BAD is the formula

$$\bigvee_{\substack{l,l'=1\\l\neq l'}}^{k}\bigvee_{i=1}^{n}(x_{li}\wedge x_{l'i})\vee\bigvee_{l=1}^{k}\bigvee_{\substack{i,j=1\\i\neq j}}^{n}(x_{li}\wedge x_{lj})\vee\bigvee_{\substack{t,t'=1\\t\neq t'}}^{k-1}\bigvee_{i=1}^{n}(y_{it}\wedge y_{it'}).$$

We show how to reduce CLIQUE_k^n to FPHP_{k-1}^k in the monotone sequent calculus. The reduction was first given in [BPR97]; here we provide proofs of correctness to check that monotonicity is preserved. The strategy will be to show that the sequents

$$\mathrm{LCLIQUE}_k^n \vdash \mathrm{LFPHP}' \tag{3.25}$$

$$\mathrm{RFPHP}' \vdash \mathrm{RCLIQUE}_k^n \tag{3.26}$$

have MLK-proofs of size polynomial in n, where LFPHP' and RFPHP' are the result of replacing the variable $p_{l,t}$ by the formula $\bigvee_{i=1}^n (x_{li} \wedge y_{it})$ in LFPHP and RFPHP respectively.

Lemma 22 *Sequent (3.25) has MLK-proofs of size polynomial in n.*

Proof: Consider the following sequence of sequents with easy MLK-proofs (the notation $A \vdash B \vdash C$ stands for the sequence $A \vdash B$, $B \vdash C$):

$$\bigwedge_{l=1}^{k} \bigvee_{i=1}^{n} x_{li} \wedge \bigwedge_{i=1}^{n} \bigvee_{t=1}^{k-1} y_{it} \vdash \bigwedge_{l=1}^{k} \left(\bigvee_{i=1}^{n} x_{li} \wedge \bigwedge_{i=1}^{n} \bigvee_{t=1}^{k-1} y_{it} \right) \vdash \bigwedge_{l=1}^{k} \bigvee_{i=1}^{n} \left(x_{li} \wedge \bigvee_{t=1}^{k-1} y_{it} \right) \vdash$$
$$\vdash \bigwedge_{l=1}^{k} \bigvee_{i=1}^{n} \bigvee_{t=1}^{k-1} (x_{li} \wedge y_{it}) \vdash \bigwedge_{l=1}^{k} \bigvee_{t=1}^{k-1} \bigvee_{i=1}^{n} (x_{li} \wedge y_{it}).$$

The first derivation follows by left weakening, left $\wedge$-introduction, and commutativity; for the second derivation use distributivity and the derivable sequent $A \wedge B \vdash A$; for the third derivation use distributivity; and for the last derivation use commutativity. Finally observe that the first formula is $\mathrm{LCLIQUE}_k^n$ and the last formula is LFPHP'_{k-1} (recall the substitution of $p_{l,t}$ by $\bigvee_{i=1}^n (x_{li} \wedge y_{it})$).
$\square$

Lemma 23 *Sequent (3.26) has MLK-proofs of size polynomial in n.*

Proof: Let us write down the full expression for RFPHP'_{k-1}:

$$\bigvee_{t=1}^{k-1} \bigvee_{\substack{l,l'=1 \\ l \neq l'}}^{k} \left[\bigvee_{i=1}^{n} (x_{li} \wedge y_{it}) \wedge \bigvee_{j=1}^{n} (x_{l'j} \wedge y_{jt}) \right] \vee \bigvee_{l=1}^{k} \bigvee_{\substack{t,t'=1 \\ t \neq t'}}^{k-1} \left[\bigvee_{i=1}^{n} (x_{li} \wedge y_{it}) \wedge \bigvee_{j=1}^{n} (x_{lj} \wedge y_{jt'}) \right].$$

By distributivity we obtain

$$\bigvee_{t=1}^{k-1} \bigvee_{\substack{l,l'=1 \\ l \neq l'}}^{k} \bigvee_{i,j=1}^{n} (x_{li} \wedge y_{it} \wedge x_{l'j} \wedge y_{jt}) \vee \bigvee_{l=1}^{k} \bigvee_{\substack{t,t'=1 \\ t \neq t'}}^{k-1} \bigvee_{i,j=1}^{n} (x_{li} \wedge y_{it} \wedge x_{lj} \wedge y_{jt'}).$$

By commutativity we obtain

$$\bigvee_{t=1}^{k-1} \bigvee_{\substack{l,l'=1 \\ l \neq l'}}^{k} \bigvee_{\substack{i,j=1 \\ i \neq j}}^{n} (x_{li} \wedge y_{ti} \wedge x_{l'j} \wedge y_{tj}) \vee \bigvee_{t=1}^{k-1} \bigvee_{\substack{l,l'=1 \\ l \neq l'}}^{k} \bigvee_{i=1}^{n} (x_{li} \wedge y_{ti} \wedge x_{l'i} \wedge y_{ti}) \vee \bigvee_{l=1}^{k} \bigvee_{\substack{t,t'=1 \\ t \neq t'}}^{k-1} \bigvee_{i,j=1}^{n} (x_{li} \wedge y_{it} \wedge x_{lj} \wedge y_{jt'}).$$

Now, working on the last two big disjuncts using $A \wedge B \vdash A$ and distinguishing on $i = j$ and $i \neq j$, we obtain

$$\bigvee_{t=1}^{k-1} \bigvee_{\substack{l,l'=1 \\ l \neq l'}}^{k} \bigvee_{\substack{i,j=1 \\ i \neq j}}^{n} (x_{li} \wedge y_{it} \wedge x_{l'j} \wedge y_{jt}) \vee \mathrm{BAD}.$$

Observe that the last formula is simply $\mathrm{RCLIQUE}_k^n$, and the proof is complete. □

Corollary 6 *The sequents* CLIQUE_k^n *have MLK-proofs of size* $n^{O(1)}$.

3.5.4 Separation Results

Let $\mathcal{P}$ and $\mathcal{P}'$ be two propositional proof systems. We say that $\mathcal{P}$ is exponentially separated from $\mathcal{P}'$ if there exists a family of tautologies with $\mathcal{P}$-proofs of size $S = S(n)$ whose shortest $\mathcal{P}'$-proofs require size $2^{S^{\Omega(1)}}$.

Recall that FPHP_n^{n+1} requires exponential-size refutations in Bounded-Depth LK and in Resolution in particular [PBI93, KPW95, Hak85]. On the other hand, CLIQUE_k^n requires exponential-size refutations in Cutting Planes [BPR97, Pud97]. Putting these together with our upper bounds, we obtain the following separation results:

Theorem 11 *MLK is exponentially separated from Resolution, Bounded-Depth LK and Cutting Planes.*

Chapter 4

Lower Bounds Beyond Resolution

4.1 Discussion on Random Restrictions

This chapter is devoted to the lower bounds on the size of proofs of the Weak Pigeonhole Principle and Random CNF Formulas in any Res(2) system. The introductory chapter of the thesis contains a wide discussion on the interest and the motivation of these results. However, we deferred the discussion about the applicability of current lower bound techniques to this chapter.

Let us remind the way the random restriction method is used in order to prove lower bounds for Random CNF formulas in Resolution. Of course, this will only be a sketch of the idea [BP96]. A restriction is a partial function $\rho :\subseteq \{x_1, \ldots, x_n\} \to \{0, 1\}$ from the set of variables to $\{0, 1\}$. We can define a probability distribution $\mathcal{R}_p$ on random restrictions as follows: Every variable is left undefined with probability p and is set to 0 or 1 with probabilities $(1-p)/2$. Each variable is treated independently of the rest. The idea of the lower bound argument is as follows: Given a alleged small Resolution refutation Π of a 3-CNF formula F, argue that a random restriction $\rho \in \mathcal{R}_p$ with appropriate p will satisfy all clauses of Π containing many literals with high probability. All the clauses of the resulting refutation $\Pi|_\rho$ of $F|_\rho$ will all have not too many literals. However, if the formula F is chosen at random according to the usual distribution, and p is not too small, one can prove that every Resolution refutation of $F|_\rho$ must contain a clause with many literals. It follows that Π could not have been small in the first place.

The most difficult part of the above argument is proving that every Resolution refutation of $F|_\rho$ must contain a large clause. Proving that a clause with many literals is satisfied with high probability is relatively straightforward. Indeed, the clause

$$x_{i_1} \vee \ldots \vee x_{i_r} \vee \neg x_{j_1} \vee \ldots \vee \neg x_{j_s}$$

is satisfied with probability at least $1 - ((1+p)/2)^{r+s}$ which is pretty close to one if $r + s$ is big.

One may like to use the above argument for Res(2) systems in which disjunctions of conjunctions of two literals (so-called 2-disjunctions) are allowed in addition to clauses. It turns out that the *large clause* lemma resists in the following form: If p is not too small, every Res(2)-refutation of $F|_\rho$ must contains a 2-disjunction with many literals. Unfortunately, it is not true anymore that ρ will almost surely satisfy every 2-disjunction with many literals. Indeed, let D be the following 2-disjunction:

$$(x_1 \wedge x_2) \vee (x_1 \wedge x_3) \vee \ldots \vee (x_1 \wedge x_n).$$

A random restriction $\rho \in \mathcal{R}_p$ will leave variable x_1 unset with probability p. Thus, $D|_\rho$ will not be satisfied with probability at least p. The point is, however, that $D|_\rho$ will almost surely become equivalent to x_1, and this is a 2-disjunction with not too many literals. Thus, we do not require that every 2-disjunction with many literals is satisfied by ρ, but rather that every 2-disjunction with many literals is either satisfied or left with a few literals after ρ is applied. This is enough for the lower bound argument to go through. Proving this property of 2-disjunctions and random restrictions is the core of our lower bound and will require a non-trivial argument. In the case of the Weak Pigeonhole Principle, the difficulty is augmented by the fact that restrictions must have some particular form.

4.2 The Weak Pigeonhole Principle

4.2.1 Definitions and Overview of the Lower Bound Proof

A k*-term* is a conjunction of up to k literals. A k*-disjunction* is an (unbounded fan-in) disjunction of k-terms. If F is a k-disjunction, a 1-term of F is also called a *free-literal*. The refutation system Res(k), defined by Krajíček [Kra00], works with k-disjunctions. We recall the three inference rules: Weakening, $\wedge$-Introduction, and Cut.

Weakening:

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

Introduction of $\wedge$:

$$\frac{A \vee (l_1 \wedge \ldots \wedge l_h) \qquad B \vee (l_{h+1} \wedge \ldots \wedge l_r)}{A \vee B \vee (l_1 \wedge \ldots \wedge l_r)}$$

Cut:

$$\frac{A \vee (l_1 \wedge \ldots \wedge l_r) \qquad B \vee \neg l_1 \vee \ldots \vee \neg l_r}{A \vee B},$$

where A and B are k-disjunctions, $1 \leq h < r \leq k$, and the l_i's are literals. As usual, if l is a literal, $\bar{l}$ denotes its negation. Observe that Res(1) coincides with Resolution with the Weakening rule. The size of a Res(k)-refutation is the number of symbols in it.

Our lower bounds will work not only for Res(2), but for any sound refutation system whose allowed formulas are 2-disjunctions, and whose rules of inference have a fixed number of premises. Such rules are said to have bounded fan-in. Note that all rules of Res(2) have fan-in at most two.

As we mentioned in the introduction, our arguments are based on random restrictions. In general terms, what we do is the following. Given an unsatisfiable CNF formula F, and an alleged small Res(2)-refutation P of F, we apply a random restriction ρ, from a suitable distribution, and we get a refutation $P|_\rho$ of $F|_\rho$. The distribution on restrictions that we choose will satisfy the following two properties:

(i) $F|_\rho$ satisfies certain expansion properties,

(ii) Every 2-disjunction in $P|_\rho$ is short (measured by the number of literals that occur).

The argument will be complete since these two conditions will be shown to be contradictory.

As a contrast with the lower bound arguments for Resolution, the most difficult part of our proof is showing that property (ii) is satisfied. The conjunctions make this task more involved. In order to overcome this, we split the restriction into two parts $\rho = \rho_1\rho_2$. Then, the main contribution is showing that every large clause in $P|_{\rho_1}$ contains many free literals. That allows us show, by a standard argument, that no large clause remains in $P|_{\rho_1\rho_2}$.

For the sake of clarity of exposition, we explain this outline again in the particular case of the Weak Pigeonhole Principle. Let $G = (U \cup V, E)$ be a bipartite graph on the sets U and V of cardinality m and n respectively, where $m > n$. The G-PHP_n^m, defined by Ben-Sasson and Wigderson [BSW01], states that there is no matching of U into V. For every edge $(u, v) \in E$, let $x_{u,v}$ be a propositional variable meaning that u is mapped to v. The principle is then formalized as the conjunction of the following set of clauses:

$$x_{u,v_1} \vee \cdots \vee x_{u,v_r} \qquad u \in U,\ N_G(u) = \{v_1, \ldots, v_r\} \tag{4.1}$$

$$\bar{x}_{u,v} \vee \bar{x}_{u',v} \qquad v \in V,\ u, u' \in N_G(v),\ u \neq u'. \tag{4.2}$$

Here, $N_G(w)$ denotes the set of neighbors of w in G. Observe that if G is the complete bipartite graph K_n^m, then G-PHP_n^m coincides with the usual pigeonhole principle PHP_n^m. It is easy to see that a lower bound for the size of Res(2)-refutations of G-PHP_n^m implies the same lower bound for the size of Res(2)-refutations of PHP_n^m.

Ben-Sasson and Wigderson proved that whenever G is expanding in a sense defined next, every Resolution refutation of G-PHP_n^m must contain a clause with many literals. We observe that this result is not unique to Resolution and holds in a more general setting. Before we state the precise result, let us recall the definition of expansion:

Definition 1 [BSW01] *Let $G = (U \cup V, E)$ be a bipartite graph where $|U| = m$, and $|V| = n$. For $U' \subset U$, the* boundary *of U, denoted by $\partial U'$, is the set of vertices in V that have exactly one neighbor in U'; that is, $\partial U' = \{v \in V : |N(v) \cap U'| = 1\}$. We say that G is (m, n, r, f)*-expanding *if every subset $U' \subseteq U$ of size at most r is such that $|\partial U'| \geq f \cdot |U'|$.*

The proof of the following statement is the same as in [BSW01] for Resolution.

Theorem 12 *Let $\mathcal{S}$ be a sound refutation system with all rules having fan-in at most two. Then, if G is (m, n, r, f)-expanding, every $\mathcal{S}$-refutation of G-PHP^m_n must contain a formula that involves at least $rf/2$ distinct literals.*

Proof: For every $u \in U$, let A_u be the conjunction of $x_{u,v_1} \vee \cdots \vee x_{u,v_s}$ and all hole axioms. Here $v_1, \ldots, v_s$ are the neighbors of u in G. For a set $U' \subseteq U$, let $A_{U'}$ be $\bigwedge_{u \in U'} A_u$. For every formula F in the refutation, let $\mu(F)$ be the minimum size of a $U' \subseteq U$ such that $A_{U'} \models F$. Obviously, $\mu(F) \leq 1$ for every axiom. On the other hand, $\mu(\Box) \geq r$ since every U' of size at most r has a matching into V. Moreover, if F is derived from F_1 and F_2, then $\mu(F) \leq \mu(F_1) + \mu(F_2)$. Here we used the soundness of the rules and that their fan-in is at most two. It follows that some F exists with $r/2 \leq \mu(F) < r$. Let $U' \subseteq U$ be minimal so that $A_{U'} \models F$.

Our goal is to prove that for every $v \in \partial U'$, there exists a $u^* \in N_G(v)$ so that $x_{u^*,v}$ occurs in F. This will justify at least $|\partial U'| \geq f \cdot |U'| \geq rf/2$ literals in F. Fix $v \in \partial U'$ and let u be its unique neighbor in U'. Let α be a truth assignment that satisfies $A_{U'-\{u\}}$ and falsifies both A_u and F. Such an assignment exists by the minimality of U'. Let α' be the result of setting $x_{u,v}$ to 1, setting $x_{u',v}$ to 0 for all $u' \in N_G(v) - \{u\}$, and leaving the rest of variables as in α. Notice that α' satisfies $A_{U'}$ since u is the unique neighbor of v in U'. It follows that α' satisfies F since $A_{U'} \models F$. Since α does not satisfy F and the only differences between α and α' are in the variables of hole v, we conclude that $x_{u^*,v}$ occurs in F for some $u^* \in N_G(v)$. $\Box$

With these definitions and results, we are ready to outline the argument of the lower bound proof. In section 4.2.2, we will prove the existence of a bipartite graph $G = (U \cup V, E)$ with $|U| = cn'$ and $|V| = n'$ such that if we remove a small random subset of nodes from V, and the corresponding edges, the resulting graph is (m, n, r, f)-expanding for certain m, n, r and f. Then we will argue that G-$\mathrm{PHP}^{cn'}_{n'}$ requires exponential-size Res(2)-refutations as follows. Assume, for contradiction, that Π is a small refutation of G-$\mathrm{PHP}^{cn'}_{n'}$. We say that a 2-disjunction in Π is large if it contains at least $d = rf/2$ distinct literals. We apply a random restriction ρ_1 to the refutation such that for every large C, either $C|_{\rho_1}$ contains many free literals, or the total number of literals in $C|_{\rho_1}$ is less than d. Then we extend ρ_1 to a new random restriction $\rho \supseteq \rho_1$ that knocks out all those large C such that $C|_{\rho_1}$ contains many free literals, ignoring those that are not free. After applying ρ, we obtain a refutation of $G(\rho)$-PHP^m_n where all 2-disjunctions have less than $rf/2$ literals and $G(\rho)$ is (m, n, r, f)-expanding. This contradicts Theorem 12.

4.2.2 Random Graphs and Restrictions

In this section we will prove the existence of a bipartite graph G as claimed in Section 4.2.1.

Let $\mathcal{G}(m,n,p)$ denote the distribution on bipartite graphs on sets U and V of sizes m and n respectively, with edge probability p independently for each edge.

Lemma 24 *If G is drawn from $\mathcal{G}(m,n,p)$, then* $\Pr\left[(\forall v \in V)(mp/2 < \deg_G(v) < 2mp)\right] \geq 1 - 2ne^{-\frac{mp}{8}}$.

Proof: Fix a vertex $v \in V$. Then, $\deg_G(v)$ is distributed as a Binomial distribution $\mathrm{Bin}(m,p)$ with parameters m and p, so that its expectation $\mathrm{E}[\deg_G(v)]$ is mp. By Chernoff bounds, $\Pr[\deg_G(v) \geq 2mp] \leq e^{-mp/3}$ and $\Pr[\deg_G(v) \leq mp/2] \leq e^{-mp/8}$. By a union bound,

$$\Pr\left[(\exists v \in V)(\deg_G(v) \leq mp/2 \vee \deg_G(v) \geq 2mp)\right] \leq ne^{-mp/3} + ne^{-mp/8} \leq 2ne^{-mp/8},$$

and so

$$\Pr\left[(\forall v \in V)(mp/2 < \deg_G(v) < 2mp)\right] \geq 1 - 2ne^{-mp/8}$$

as required. □

Lemma 25 *Let $m = kn$, $p = 48k\ln(m)/m$, $\alpha = 1/mp$ and $f = np/6$. Let G be drawn from $\mathcal{G}(m,n,p)$. Then,* $\Pr[G$ *is $(m,n,\alpha m,f)$-expanding*$] \geq 1/2$.

Proof: Fix $U' \subseteq U$ of size $s \leq \alpha m$, and $v \in V$. Then,

$$\Pr[v \in \partial U'] = sp(1-p)^{s-1}.$$

Let $q = \Pr[v \in \partial U']$. Let X_v be the indicator random variable for the event that $v \in \partial U'$. Then, $|\partial U'| = \sum_{v \in V} X_v$. Observe that X_v and $X_{v'}$ are independent whenever $v \neq v'$. Hence, $|\partial U'| \sim \mathrm{Bin}(n,q)$, so that $\mathrm{E}[|\partial U'|] = nq$. By Chernoff bound, $\Pr[|\partial U'| \leq nq/2] \leq e^{-nq/8}$. On the other hand,

$$nq = nsp(1-p)^{s-1} \geq snp(1-p)^{\alpha m}.$$

Moreover, $(1-p)^{\alpha m} = (1-p)^{1/p}$ approaches $1/e$ for sufficiently large m. Therefore, $nq \geq snp/3$. It follows that $nq/2 \geq sf$ and $e^{-nq/8} \leq e^{-snp/24}$. We conclude that

$$\Pr[|\partial U'| < f \cdot |U'|] \leq \Pr[|\partial U'| \leq nq/2] \leq e^{-nq/8} \leq e^{-snp/24}.$$

Finally, we bound the probability that G is not $(m,n,\alpha m,f)$-expanding by

$$\sum_{s=1}^{\alpha m} \binom{m}{s} e^{-snp/24} \leq \sum_{s=1}^{\alpha m} m^s e^{-snp/24} \leq \sum_{s=1}^{\alpha m} (me^{-np/24})^s. \tag{4.3}$$

Recall that $p = 48k\ln(m)/m$ and $m = kn$. So $me^{-np/24} \leq me^{-2\ln(m)} = m^{-1} < 1/4$. Hence the sum in (4.3) is bounded by

$$\sum_{s=1}^{\infty} \frac{1}{4^s} \leq \frac{1}{2}.$$

□

Let G be a fixed bipartite graph on $\{1, \ldots, m\}$ and $\{1, \ldots, n\}$. A *restriction* (for G) is a sequence of pairs $\rho = ((u_1, v_1), \ldots, (u_r, v_r))$ such that $(u_i, v_i) \in E(G)$, and all v_i's are distinct. We let $R_r(G)$ be the set of restrictions of length r. We define a distribution $\mathcal{R}_r(G)$ on $R_r(G)$ as follows: Let $V_0 = \{1, \ldots, n\}$; for every $i \in \{1, \ldots, r\}$ in increasing order, choose a hole v_i uniformly at random in V_{i-1}, choose a pigeon u_i uniformly at random in $N_G(v_i)$, and let $V_i = V_{i-1} - \{v_i\}$. The final restriction is $((u_1, v_1), \ldots, (u_r, v_r))$.

We define a distribution $\mathcal{D}(m, n, p, r)$ on the set of pairs (G, ρ) with $\rho \in R_r(G)$: the graph G is drawn from $\mathcal{G}(m, n + r, p)$ first, and then ρ is drawn from $\mathcal{R}_r(G)$. In other words, if (H, π) is a fixed pair with $\pi \in R_r(H)$, then

$$\Pr[G = H \wedge \rho = \pi] = \frac{p^{e(H)}(1-p)^{m(n+r)-e(H)}}{|R_r(H)|}.$$

If G is a bipartite graph on the vertex sets $\{1, \ldots, m\}$ and $\{1, \ldots, n + r\}$, and ρ is a restriction $((u_1, v_1), \ldots, (u_r, v_r)) \in R_r(G)$, then $G(\rho)$ denotes the graph that results from deleting $v_1, \ldots, v_r$ from G, and renaming nodes in an order-preserving way. With this definitions we are ready to prove:

Lemma 26 *Let $m = kn$, $p = 48k\ln(m)/m$, $\alpha = 1/mp$ and $f = np/6$. Let (G, ρ) be drawn from $\mathcal{D}(m, n, p, r)$. Then, $\Pr[G(\rho)$ is $(m, n, \alpha m, f)$-expanding$] \geq 1/2$.*

Proof: Let A be the event that $G(\rho)$ is $(m, n, \alpha m, f)$-expanding. Let

$$S = \{R \subseteq \{1, \ldots, n + r\} : |R| = r\}.$$

Then,

$$\Pr[A] = \sum_{R \in S} \Pr[A \mid \mathrm{ran}(\rho) = R] \Pr[\mathrm{ran}(\rho) = R].$$

The proof that $\Pr[A \mid \mathrm{ran}(\rho) = R] \geq 1/2$ is the same as the proof of Lemma 25 replacing V by $V - R$. The result follows. □

Lemma 27 *Let $m = kn$, $p = 48k\ln(m)/m$, $\alpha = 1/mp$ and $f = np/6$. For every $r \leq n$, there exists a bipartite graph H on $\{1, \ldots, m\}$ and $\{1, \ldots, n + r\}$ such that the following two properties hold:*

(i) $mp/2 \leq \deg_H(v) \leq 2mp$ *for every* $v \in \{1, \ldots, n + r\}$,

(ii) $\Pr[H(\rho)$ *is* $(m, n, \alpha m, f)$*-expanding*$] \geq 1/3$,

when ρ is drawn from $\mathcal{R}_r(H)$.

Proof: Let (G, ρ) be drawn from $\mathcal{D}(m, n, p, r)$. We have

$$\Pr\left[G(\rho) \text{ is } (m, n, \alpha m, f)\text{-expanding}\right] \geq 1/2$$

by Lemma 26. Moreover,

$$\Pr\left[(\forall v \in V)(mp/2 < \deg_G(v) < 2mp)\right] \geq 1 - (n + r)e^{-mp/9} \geq 5/6$$

by Lemma 24. Let $E(G, \rho)$ be the event that $G(\rho)$ is expanding *and* every right-node in G has degree between $mp/2$ and $2mp$. Combining both equations above we have that $\Pr\left[E(G, \rho)\right] \geq 1/3$. On the other hand,

$$\Pr\left[E(G, \rho)\right] = \sum_H \Pr\left[E(G, \rho) \mid G = H\right] \Pr\left[G = H\right]$$

where H ranges over all bipartite graphs on m and $n + r$ nodes. Therefore, there exists some fixed H such that

$$\Pr\left[E(G, \rho) \mid G = H\right] \geq 1/3.$$

Moreover, $\Pr\left[E(G, \rho) \mid G = H\right]$ equals $\Pr\left[E(H, \pi)\right]$ when π is drawn from $\mathcal{R}_r(H)$. Finally, since this probability is strictly positive, it must be the case that H satisfies property (i) in the lemma since it is independent of π. □

4.2.3 The Lower Bound Argument

Before we state and prove our main theorem, we will give some definitions and lemmas.

Let us first give a normal form for Res(2)-refutations of G-PHP^m_n. We claim that every Res(2)-refutation of G-PHP^m_n can be turned into a Res(2)-refutation of similar size in which no 2-term is of the form $x_{u,v} \wedge x_{u',v}$ with $u \neq u'$. To check this, observe that such a 2-term must have been introduced at some point by the rule of $\wedge$-introduction with, say, $A \vee x_{u,v}$ and $B \vee x_{u',v}$. Cutting them with the axiom $\bar{x}_{u,v} \vee \bar{x}_{u',v}$ we get $A \vee B$ that can be used to continue the proof because it subsumes $A \vee B \vee (x_{u,v} \wedge x_{u',v})$.

Let C be a 2-disjunction, and let $(u, v) \in E(G)$. We let $C|_{(u,v)}$ be the result of assigning $x_{u,v} = 1$ and $x_{u',v} = 0$ for every $u' \in N_G(v) - \{u\}$ to C, and simplifying as much as possible. This includes replacing subformulas of the form $l \vee (l \wedge l')$ by l in some specified order; here l and l' are literals. Given a restriction $\rho = ((u_1, v_1), \ldots, (u_r, v_r))$, we let $C|_\rho$ be the result of applying $(u_1, v_1), \ldots, (u_r, v_r)$ successively in this order. For every $i \in \{1, \ldots, r\}$, we let $\rho_i = ((u_1, v_1), \ldots, (u_i, v_i))$.

Let us now study in more detail the result of applying a pair of a restriction to a 2-disjunction. First we give some definitions. We say that a pair $(u, v) \in E(G)$ *hits* C if either $x_{u,v}$ occurs positively in C, or $x_{u',v}$ occurs negatively in C for some $u' \in N_G(v) - \{u\}$. Equivalently, (u, v)

hits C if it sets some literal of C to 1. If the literal is free, it knocks out the 2-disjunction. If the literal is part of a conjunction, it will locally create a free literal. In general, we say that $(u, v) \in E(G)$ *knocks* C if $C|_{(u,v)} \equiv 1$. We say that $(u, v) \in E(G)$ is a *bad choice* for C if it does not knock it and there exists $u' \in N_G(v) - \{u\}$ such that (u', v) knocks C. A bad choice may or may not be a hit.

Lemma 28 *Let C be a simplified 2-disjunction, and $(u, v) \in E(G)$. If (u, v) hits C and is not a knock or a bad choice, then $C|_{(u,v)}$ has more free literals than C.*

Proof: First notice that the literals that (u, v) sets to 1 are in a conjunction, otherwise (u, v) is a knock. Such literals can appear positively or negatively. We will discuss the two cases:

Case 1: The literal is $x_{u,v}$ and appears in a conjunction of the form $x_{u,v} \wedge y$. The pair (u, v) does not set y to 1 otherwise we would have a knock. Also, it does not set it to 0 either, otherwise $y = x_{u',v}$ and such a conjunction is not allowed in the normal form. On the other hand, y does not appear free because C is a simplified 2-disjunction. Finally no free literal of C desapears when we apply (u, v) to C, otherwise (u, v) would be a bad choice.

Case 2: The literal is $\overline{x}_{u',v}$, and it appears in a conjunction of the form $\overline{x}_{u',v} \wedge y$. Because (u, v) is not a knock, it does not set y to 1. Also, (u, v) does not set y to 0 either, otherwise it would be a bad choice, given that the indegree of v is 3 or more. As in the previous case and for the same reasons, y does not appear free in C, and no free literal of C disappears when we apply (u, v). $\square$

Theorem 13 *Let $c > 1$ be a constant. For all sufficiently large n, every Res(2)-refutation of* PHP^{cn}_n *has size at least $e^{n/(\log n)^{15}}$.*

Proof: Let $k = c+1$, $r = n/c$, $n' = n+r$, and $m = kn = cn'$. Let $G = (U \cup V, E)$ with $|U| = m$ and $|V| = n + r$ be the bipartite graph of Lemma 27. We show that every Res(2)-refutation of G-$\mathrm{PHP}^{cn'}_{n'}$ has size at least $e^{n/(\log n)^{14}}$. This will imply the Theorem since a Res(2)-refutation of $\mathrm{PHP}^{cn'}_{n'}$ gives a Res(2)-refutation of G-$\mathrm{PHP}^{cn'}_{n'}$ of no bigger size. Let us assume, for contradiction, that G-$\mathrm{PHP}^{cn'}_{n'}$ has a Res(2)-refutation Π of size $S < e^{n/(\log n)^{14}}$.

We will use the following concepts. We say that C is *large* if it contains at least $d = n/12$ distinct literals; otherwise, C is *small*. We say that C is *wide* if it contains at least $s = n/(\log n)^5$ free literals; otherwise, C is *narrow*.

In all probabilities that follow, ρ is drawn from the distribution $\mathcal{R}_r(G)$. Our main goal is to prove that the probability that a fixed 2-disjunction C of Π remains large is exponentially small; that is, we aim for a proof that

$$\Pr\left[C|_\rho \text{ is large}\right] \leq e^{-n/(\log n)^{13}}. \tag{4.4}$$

This will suffice because then

$$\Pr\left[(\exists C \in \Pi)(C|_\rho \text{ is large})\right] \leq S e^{-n/(\log n)^{13}} < 1/3,$$

and also

$$\Pr\left[G(\rho) \text{ not } (m, n, \alpha m, f)\text{-expanding}\right] \leq 2/3$$

by Lemma 27. Combining, this means that there exists a restriction $\rho \in R_r(G)$ such that $G(\rho)$ is $(m, n, \alpha m, f)$-expanding and every 2-disjunction in $\Pi|_\rho$ has less than $d = \alpha m f/2$ literals. This is a contradiction with Theorem 12 since $\Pi|_\rho$ is a sound refutation.

For $i \in \{1, \ldots, r\}$, let A_i be the event that $C|_{\rho_i}$ is large, and let B_i be the event that $C|_{\rho_i}$ is narrow. Recall that $\rho_i = ((u_1, v_1), \ldots, (u_i, v_i))$. Then,

$$\begin{aligned}
\Pr\left[C|_\rho \text{ is large}\right] &\leq \Pr\left[A_r \wedge \bigvee_{j \geq r/2} B_j\right] + \Pr\left[A_r \wedge \bigwedge_{j \geq r/2} \overline{B_j}\right] \leq \\
&\leq \sum_{j=r/2}^{r} \Pr\left[A_j \wedge B_j\right] + \Pr\left[A_r \wedge \bigwedge_{j \geq r/2} \overline{B_j}\right].
\end{aligned}$$

We will show that every term in this expression is exponentially small. The bound on terms of the form $\Pr\left[A_j \wedge B_j\right]$ will be proven in Lemma 30. For the last term, we use an argument similar in spirit to the one by Beame and Pitassi [BP96]:

Lemma 29 $\Pr\left[A_r \wedge \bigwedge_{j \geq r/2} \overline{B_j}\right] \leq e^{-n/(\log n)^8}$.

Proof: Let S_i be the indicator random variable for the event that (u_i, v_i) knocks $C|_{\rho_{i-1}}$. Then,

$$\begin{aligned}
\Pr\left[A_r \wedge \bigwedge_{j \geq r/2} \overline{B_j}\right] &\leq \Pr\left[\bigwedge_{i > r/2} S_i = 0 \wedge \bigwedge_{j \geq r/2} \overline{B_j}\right] = \\
&= \prod_{i > r/2} \Pr\left[S_i = 0 \wedge \bigwedge_{j \geq r/2} \overline{B_j} \,\middle|\, \bigwedge_{r/2 < j < i} S_j = 0\right] \leq \\
&\leq \prod_{i > r/2} \Pr\left[S_i = 0 \wedge \overline{B_{i-1}} \,\middle|\, \bigwedge_{r/2 < j < i} S_j = 0\right] \leq \\
&\leq \prod_{i > r/2} \Pr\left[S_i = 0 \,\middle|\, \overline{B_{i-1}} \wedge \bigwedge_{r/2 < j < i} S_j = 0\right].
\end{aligned}$$

Fix $i \in \{r/2 + 1, \ldots, r\}$ and let H be the set of holes that occur in a free literal of $C|_{\rho_{i-1}}$. Given that $\overline{B_{i-1}}$ holds, $C|_{\rho_{i-1}}$ is wide which means that there are at least s free literals. Therefore $|H| \geq s/2\Delta$, where $\Delta = 2mp$ is an upper bound on the right-degree of G. Moreover, every $v \in H$ gives a possible knock, and different holes give different knocks. The reason is the following: if $x_{u,v}$ is a free literal, then (u, v) is a knock; and if $\bar{x}_{u,v}$ is a free literal, then (u', v) is a knock for every $u' \in N_G(v) - \{u\}$, which is non-empty since the right-degree of G is at least two. Therefore,

$$\Pr\left[S_i = 1 \,\middle|\, \overline{B_{i-1}} \wedge \bigwedge_{r/2 < j < i} S_j = 0\right] \geq \frac{|H|}{\Delta(n + r - i + 1)} \geq \frac{s}{3\Delta^2 n}.$$

Therefore,

$$\Pr\left[A_r \wedge \bigwedge_{j \geq r/2} \overline{B_j}\right] \leq \left(1 - \frac{s}{3\Delta^2 n}\right)^{r/2} \leq e^{-\frac{sr}{6\Delta^2 n}} \leq e^{-n/(\log n)^8}.$$

□

Lemma 30 *Let j be such that $r/2 \leq j \leq r$. Then,* $\Pr[A_j \wedge B_j] \leq e^{-n/(\log n)^{11}}$.

Proof: Recall that A_j is the event that $C|_{\rho_j}$ is large, and B_j is the event that $C|_{\rho_j}$ is narrow. We let S_i be the indicator random variable for the event that (u_i, v_i) hits $C|_{\rho_{i-1}}$, where $\rho_{i-1} = ((u_1, v_1), \ldots, (u_{i-1}, v_{i-1}))$. Let $S = \sum_{i=1}^{j} S_i$. Then, for every h,

$$\begin{aligned} \Pr[A_j \wedge B_j] &= \Pr[A_j \wedge B_j \wedge S < h] + \Pr[A_j \wedge B_j \wedge S \geq h] \leq \\ &\leq \Pr[A_j \wedge S < h] + \Pr[A_j \wedge B_j \wedge S \geq h]. \end{aligned}$$

We show that each term in this expression is exponentially small. More precisely, we show that $\Pr[A_j \wedge S < h] \leq e^{-n/(\log n)^3}$ and $\Pr[A_j \wedge B_j \wedge S \geq h] \leq e^{-n/(\log n)^{10}}$ which is clearly enough to prove Lemma 30.

Claim 1 *Let $h = n/(\log n)^4$. Then,* $\Pr[A_j \wedge S < h] \leq e^{-n/(\log n)^3}$.

Proof: Let $Y = \{(a_1, \ldots, a_j) \in \{0,1\}^j : \sum_{i=1}^{j} a_i < h\}$. Observe that A_j implies A_i for every $i \leq j$ because if $C|_{\rho_j}$ is large, so is $C|_{\rho_i}$ for every $i \leq j$. Then,

$$\begin{aligned} \Pr[A_j \wedge S < h] &= \Pr\left[\sum_{i=1}^{j} S_i < h \wedge A_j\right] = \sum_{\overline{a} \in Y} \Pr\left[\bigwedge_{i=1}^{j} S_i = a_i \wedge A_j\right] = \\ &= \sum_{\overline{a} \in Y} \prod_{i=1}^{j} \Pr\left[S_i = a_i \wedge A_j \;\middle|\; \bigwedge_{k=1}^{i-1} S_k = a_k\right] \leq \\ &\leq \sum_{\overline{a} \in Y} \prod_{i=1}^{j} \Pr\left[S_i = a_i \wedge A_{i-1} \;\middle|\; \bigwedge_{k=1}^{i-1} S_k = a_k\right] \leq \\ &\leq \sum_{\overline{a} \in Y} \prod_{i=1}^{j} \Pr\left[S_i = a_i \;\middle|\; A_{i-1} \wedge \bigwedge_{k=1}^{i-1} S_k = a_k\right]. \end{aligned}$$

Fix $i \in \{1, \ldots, j\}$. Let H be the set of holes that occur in $C|_{\rho_{i-1}}$. We have $|H| \geq d/2\Delta$ given that A_{i-1} holds. Again, $\Delta = 2mp$ is an upper bound to the right-degree of G. Moreover, every $v \in H$ gives a possible hit, and different holes give different hits (the reason is the same as in Lemma 29 for knocks). Therefore,

$$\Pr\left[S_i = 1 \;\middle|\; A_{i-1} \wedge \bigwedge_{k=1}^{i-1} S_k = a_k\right] \geq \frac{|H|}{\Delta(n + r - i + 1)} \geq \frac{d}{3\Delta^2 n}.$$

Since there are at least $j-h$ zeros in $(a_1,\ldots,a_j)$, we obtain

$$\begin{aligned}\Pr\left[A_j \wedge S < h\right] &\le \sum_{\overline{a}\in Y}\left(1-\frac{d}{3\Delta^2 n}\right)^{j-h} \le \sum_{i<h}\binom{j}{i} e^{-\frac{d(j-h)}{3\Delta^2 n}} \le h j^h e^{-\frac{d(j-h)}{3\Delta^2 n}} \le \\ &\le \exp\left(-\frac{j-h}{36\Delta^2}+h\log(j)+\log(h)\right) \le e^{-n/(\log n)^3}.\end{aligned}$$

□ (of claim 1)

Claim 2 $\Pr\left[A_j \wedge B_j \wedge S \ge h\right] \le e^{-n/(\log n)^{10}}$.

Proof: During this proof we will drop the subindex j in A_j and B_j since it will always be the same. For every $i \in \{1,\ldots,r\}$, let $T_i \in \{\mathrm{k},\mathrm{b},\mathrm{n}\}$ be a random variable indicating whether (u_i,v_i) is a knock, a bad choice, or none of the previous respectively for $C|_{\rho_{i-1}}$. For $t \in \{\mathrm{k},\mathrm{b},\mathrm{n}\}$, let S_i^t be the indicator random variable for the event that $T_i = t$, and let $S^t = \sum_{i=1}^{j} S_i^t$. Thus, S^{k} is the number of knocks and S^{b} is the number of bad choices of ρ_j.

Fix $\rho = ((u_1,v_1),\ldots,(u_r,v_r))$ such that $A \wedge B \wedge S \ge h$ holds under ρ. Observe that (u_i,v_i) does not knock $C|_{\rho_{i-1}}$ for any $i \in \{1,\ldots,j\}$ since $C|_{\rho_j}$ must be large. Hence, $S^{\mathrm{k}} = 0$ under ρ. Let $b = (h-s)/(2\Delta+1)$. We now claim that $S^{\mathrm{b}} \ge b$. Suppose for contradiction that the number of bad choices is less than b. Every bad choice (u_i,v_i) removes at most 2Δ free literals since at most those many literals about hole v_i may appear. Moreover, since there are no knocks, every hit (u_i,v_i) that is not a bad choice increases the number of free literals by at least one (see lemma 28). It follows that the number of free literals in $C|_{\rho_j}$ is at least

$$(S-S^{\mathrm{b}}) - 2\Delta S^{\mathrm{b}} > h-(2\Delta+1)b = s,$$

a contradiction with the fact that B holds under ρ. We have proved that

$$\Pr\left[A \wedge B \wedge S \ge h\right] \le \Pr\left[S^{\mathrm{k}} = 0 \wedge S^{\mathrm{b}} \ge b\right].$$

The intuition behind why this last probability is small is that every bad choice could have been a knock. This makes unlikely that ρ produces many bad choices and no knocks. In what follows, we will prove this intuition using martingales.

For $t \in \{\mathrm{k},\mathrm{b},\mathrm{n}\}$ and $i \in \{1,\ldots,j\}$, let P_i^t be the random variable $\Pr\left[T_i = t \mid \rho_0,\ldots,\rho_{i-1}\right]$. We define a martingale $X_0,\ldots,X_j$ with respect to $\rho_0,\ldots,\rho_j$ as follows: Let

$$\begin{aligned}X_0 &= 0,\\ X_{i+1} &= X_i + S_{i+1}^{\mathrm{b}} - P_{i+1}^{\mathrm{b}}.\end{aligned}$$

Recall that S_{i+1}^{b} is the indicator random variable for the event that $T_{i+1} = \mathrm{b}$. So

$$\begin{aligned}\mathrm{E}[X_{i+1} \mid \rho_0,\ldots,\rho_i] &= (X_i + 1 - P_{i+1}^{\mathrm{b}})\cdot P_{i+1}^{\mathrm{b}} + (X_i - P_{i+1}^{\mathrm{b}})\cdot(1-P_{i+1}^{\mathrm{b}}) = \\ &= (X_i - P_{i+1}^{\mathrm{b}})(P_{i+1}^{\mathrm{b}} + 1 - P_{i+1}^{\mathrm{b}}) + P_{i+1}^{\mathrm{b}} = X_i.\end{aligned}$$

Hence, $\{X_i\}_i$ is a martingale with respect to $\{\rho_i\}_i$. Observe also that $X_j = S^{\rm b} - \sum_{i=1}^{j} P_i^{\rm b}$. Similarly, we define $Y_0, \ldots, Y_j$ as follows: Let

$$\begin{aligned} Y_0 &= 0, \\ Y_{i+1} &= Y_i + S_{i+1}^{\rm k} - P_{i+1}^{\rm k}. \end{aligned}$$

It is also easy to see that $\{Y_i\}_i$ is a martingale with respect to $\{\rho_i\}_i$. Again, $Y_j = S^{\rm k} - \sum_{i=1}^{j} P_i^{\rm k}$.

Subclaim 1 $P_i^{\rm k}(\rho) \geq P_i^{\rm b}(\rho)/\Delta$ *for every* $\rho \in R_r(G)$ *and* $i \in \{1, \ldots, j\}$.

Proof: Fix $i \in \{1, \ldots, j\}$ and $\rho = ((u_1, v_1), \ldots, (u_r, v_r))$. We want to show that $P_i^{\rm k}(\rho) \geq P_i^{\rm b}(\rho)/\Delta$. Define three sets as follows: let $Q = \{(u, v) \in E(G) : v \notin \{v_1, \ldots, v_{i-1}\}\}$, let $Q^{\rm k}$ be the set of knocks for $C|_{\rho_{i-1}}$ in Q, and let $Q^{\rm b}$ be the set of bad choices for $C|_{\rho_{i-1}}$ in Q. Observe that $P_i^{\rm b}(\rho) = |Q^{\rm b}| \cdot |Q|^{-1}$ and $P_i^{\rm k}(\rho) = |Q^{\rm k}| \cdot |Q|^{-1}$. On the other hand, every bad choice $(u, v) \in Q^{\rm b}$ gives a possible knock $(u', v) \in Q^{\rm k}$ by definition. Moreover, bad choices with different hole components give different possible knocks. Grouping $Q^{\rm b}$ by holes, we have that $|Q^{\rm k}| \geq |Q^{\rm b}|/\Delta$. Consequently, $P_i^{\rm k}(\rho) \geq P_i^{\rm b}(\rho)/\Delta$ as required. □ (of subclaim 1)

To complete the proof of claim 2 we will need the following form of Azuma's Inequality: Let $X_0, \ldots, X_n$ be a martingale such that $|X_i - X_{i-1}| \leq 1$; then, $\Pr[|X_n - X_0| \geq \lambda] \leq 2e^{-\lambda^2/n}$ for every $\lambda > 0$ [GS82]. Now,

$$\begin{aligned} \Pr[S^{\rm k} = 0 \wedge S^{\rm b} \geq b] &= \Pr[S^{\rm k} = 0 \wedge S^{\rm b} \geq b \wedge X_j \geq b/2] + \\ &+ \Pr[S^{\rm k} = 0 \wedge S^{\rm b} \geq b \wedge X_j < b/2]. \end{aligned}$$

The first summand is bounded by $\Pr[X_j \geq b/2] \leq 2e^{-b^2/4j}$ by Azuma's Inequality. The second summand is bounded by

$$\begin{aligned} \Pr\left[S^{\rm k} = 0 \wedge \textstyle\sum_{i=1}^{j} P_i^{\rm b} \geq b/2\right] &\leq \Pr\left[S^{\rm k} = 0 \wedge \textstyle\sum_{i=1}^{j} P_i^{\rm k} \geq b/2\Delta\right] \leq \\ &\leq \Pr[Y_j \leq -b/2\Delta] \leq 2e^{-b^2/4\Delta^2 j}. \end{aligned}$$

The first inequality follows from Subclaim 1, and the third follows from Azuma's Inequality again. The addition of the two summands is then bounded by $e^{-n/(\log n)^{10}}$ as required. □ (of claim 2 and lemma 30)

We are ready to complete the proof of our goal: equation (4.4). We have shown that

$$\Pr[C|_\rho \text{ large}] \leq \sum_{j=r/2}^{r} e^{-n/(\log n)^{11}} + e^{-n/(\log n)^{8}} \leq e^{-n/(\log n)^{13}}.$$

□

By a different setting of parameters, it is easy to see that the strongest lower bound for PHP_n^m is of the form $e^{\frac{n^9}{(m \log m)^8 \log^3 n}}$. Namely, put $r = n/8$, $h = n^3/((m \log m)^2 \log n)$ and $s = h/2$ for that calculation. Therefore the best result is an exponential lower bound for $\mathrm{PHP}_n^{n^{9/8-\epsilon}}$.

We conclude this section with a separation result. Given that Res($\log$) and depth-0.5 LK are polynomially equivalent, and given that PHP_n^{2n} has quasipolynomial-size proofs in depth-0.5 LK [MPW00], we obtain:

Corollary 7 *There is an exponential separation between Res(2) and Res($\log$).*

We note that this result indicates that any possible extension of our technique to Res(k) for $k > 2$ would have to deal with the fact that for relatively small k (polylog in n), the lower bounds cannot be exponential. Although this is valuable information, it reveals that the analysis must be quite tight.

4.3 Random Formulas

4.3.1 Random Formulas and Restrictions

The model of random k-CNF formulas that we use is the one considered in [CS88, BKPS98]. The distribution is denoted $\mathcal{F}_m^{k,n}$ and consists in choosing m clauses of exactly k literals independently with replacement. Most of the next definitions are taken and adapted from [BKPS98].

Definition 2 *For a real number n, a set of clauses $\mathcal{C}$ is n-*sparse *if $|\mathcal{C}| \leq n|v(\mathcal{C})|$ where $v(\mathcal{C})$ is the set of variables appearing in $\mathcal{C}$.*

Definition 3 *If $\mathcal{C}$ is a set of clauses and l is a literal, we say that l is* pure *in $\mathcal{C}$ if some clause of $\mathcal{C}$ contains l and no clause of $\mathcal{C}$ contains $\bar{l}$.*

Definition 4 *For $s \leq 1$ and $\epsilon \in (0, 1)$, the following properties are defined for CNF formulas F:*

- *$A(s)$: Every set of $r \leq s$ clauses of F is 1-sparse.*
- *$B_\epsilon(s)$: For r such that $s/2 < r \leq s$, every subset of r clauses of F has at least ϵr pure literals.*

For a given refutation system $\mathcal{S}$, we say that an $\mathcal{S}$-refutation is k-bounded if all formulas of the refutation involve at most k distinct literals.

Proposition 3 [BKPS98] *Let $\mathcal{S}$ be a sound refutation system with all rules of fan-in at most two. Let $s > 0$ be an integer and F be a CNF formula. If properties $A(s)$ and $B_\epsilon(s)$ both hold for F, then F has no $\epsilon s/2$-bounded $\mathcal{S}$-refutation.*

A restriction is a sequence of pairs (x, v) where x is a variable and v is either $true$ or $false$. For a 2-disjunction C let $|C|$ be the number of distinct literals occurring in it. Let $\mathcal{R}$ be a probability distribution on restrictions. We say that $\mathcal{R}$ satisfies property $R(d, M)$ if and only if for every 2-disjunction C, $\Pr\left[|C|_\rho| \geq d\right] \leq 1/M$. We will consider two probability distributions.

- $\mathcal{A}_t$ chooses a permutation of the variables uniformly at random, then chooses each variable with probability t/n in the order of the permutation. The values assigned to the variables are chosen uniformly at random from $true$ and $false$.

- $\mathcal{B}_t$ chooses r, the length of the restriction, with a binomial distribution of parameters t/n and n. Then chooses uniformly at random any sequence of variables of length r without repetitions. The values assigned to the variables are chosen uniformly at random from $true$ and $false$.

We prove that $\mathcal{A}_t$ and $\mathcal{B}_t$ are the same probability distribution. Obviously both experiments produce exactly the same restriccions. We only need to show that every restriction ρ has the same probability in both spaces.

Lemma 31 *For every $x_1, \ldots, x_r$ and $v_1, \ldots, v_r$,*

$$\Pr_{\rho\sim\mathcal{A}_t}[\rho = ((x_1, v_1), \ldots, (x_r, v_r))] = \Pr_{\rho\sim\mathcal{B}_t}[\rho = ((x_1, v_1), \ldots, (x_r, v_r))].$$

Proof: The probability $\Pr_{\rho\sim\mathcal{B}_t}[\rho = ((x_1, v_1), \ldots, (x_r, v_r))]$ is easy:

$$\binom{n}{r}\left(\frac{t}{n}\right)^r\left(1-\frac{t}{n}\right)^{n-r}\frac{1}{n2(n-1)2\ldots(n-r+1)2}. \tag{4.5}$$

The first part corresponds to the probability of choosing the value r from a binomial distribution. Remember that r is the length of the restriction. The rest of the expression is the probability of choosing the r correct pairs (x_i, v_i).

The probability $\Pr_{\rho\sim\mathcal{A}_t}[\rho = ((x_1, v_1), \ldots, (x_r, v_r))]$ is a little trickier. We will compute the probability of finding a permutation of the variables that is compatible with $(x_1, \ldots, x_r)$, that is, the variables $\{x_1, \ldots, x_r\}$ appear in that order. Then we multiply this probability by the probability of choosing the exact places where the variables in ρ are and choosing the right value for them:

$$\frac{\binom{n}{r}(n-r)!}{n!}\left(\frac{t}{n}\right)^r\left(1-\frac{t}{n}\right)^{n-r}\frac{1}{2^r}. \tag{4.6}$$

We first choose r places to put the variables in ρ, then we fill the gaps with the permutations of the other $n - r$ variables. These are the favorable cases, those that are compatible. Straightforward manipulations show that (4.5) and (4.6) are equal. □

The following is adapted from [BKPS98], with a minor change in the probability distribution.

Lemma 32 *For each integer $k \geq 3$ and $\epsilon > 0$, there are constants c_k, $c_{k,\epsilon}$, such that the following holds. Let m, n, s, t with $m = \Delta n$ for $\Delta \geq 1$. Let $F \sim \mathcal{F}_m^{k,n}$ and $\rho \sim \mathcal{A}_t$.*

- *If $t \leq c_k n/m^{1/k}$ and $s \leq c_k n/\Delta^{1/(k-2)}$, then $F|_\rho$ satisfies $A(s)$ with probability $1 - o(1)$ in s.*
- *If $s, t \leq c_{k,\epsilon} n/\Delta^{2/(k-2-\epsilon)}$, then $F|_\rho$ satisfies $B_\epsilon(s)$ with probability $1 - o(1)$ in s.*

Theorem 14 *Let $\mathcal{F}$ be a distribution over k-CNF formulas. Let $s, M \geq 1$ and $\epsilon > 0$ and let $\mathcal{R}$ be a distribution over restrictions that satisfies $R(\epsilon s/2, M)$. Then,*

$$\begin{aligned} \Pr_{F \sim \mathcal{F}}[res2(F) < M/2] &\leq 2 \Pr_{F \sim \mathcal{F}, \rho \sim \mathcal{R}}[F|_\rho \textit{ does not satisfy } A(s)] + \\ &+ 2 \Pr_{F \sim \mathcal{F}, \rho \sim \mathcal{R}}[F|_\rho \textit{ does not satisfy } B_\epsilon(s)], \end{aligned}$$

where $res2(F)$ is the minimum size of a $\mathrm{Res}(2)$-refutation of F.

Proof: For a fixed unsatisfiable k-CNF F, let P be a minimal-size $\mathrm{Res}(2)$-refutation of F. Let $\rho \sim \mathcal{R}$.

$$\begin{aligned} \Pr[F|_\rho \text{ satisfies } A(s) \wedge B_\epsilon(s)] &\leq \Pr[P|_\rho \text{ is not } \epsilon s/2 \text{ bounded}] \\ &\leq \Pr[(\exists C \in P)(|C|_\rho| > \epsilon s/2)] \\ &\leq res2(F)\frac{1}{M}. \end{aligned}$$

The first inequality follows by Proposition 3, the second is immediate, and the third follows by union bound and the fact that $\mathcal{R}$ satisfies $R(\epsilon s/2, M)$.

To finish, let

$$p(F) = \Pr_\rho[F|_\rho \text{ does not satisfy } A(s)] + \Pr_\rho[F|_\rho \text{ does not satisfy } B_\epsilon(s)].$$

Then, $res2(F) < M/2$ implies that $\Pr_\rho[F|_\rho \text{ satisfies } A(s) \wedge B_\epsilon(s)] < 1/2$, and so $p(F) > 1/2$. Therefore, $\Pr_F[res2(F) < M/2] \leq \Pr_F[p(F) > 1/2] \leq 2E_F[p(F)]$ by Markov's inequality. The result follows. □

4.3.2 The Lower Bound Argument

For simplicity, we only state the lower bound for the case $\mathcal{F}_{5n}^{3,n}$.

Theorem 15 *Let $F \sim \mathcal{F}_{5n}^{3,n}$. Then, $\mathrm{Res}(2)$-refutations of F require size $2^{\Omega(n^{1/3}/(\log(n))^2)}$ almost surely.*

Proof: Let $m = 5n$, $k = 3$, fix an arbitrary $\epsilon \in (0, 1)$, and put $t = c_3 n/(5n)^{1/3} = c' n^{2/3}$ and $s = \min(c_3 n/5, c_{3,\epsilon} n/5^{2/1-\epsilon})$. Observe that these numbers satisfy the two hypothesis in Lemma 32. Let $M = 2^{n^{1/3}/(\log(n))^3}$. If we could prove that $\mathcal{B}_t$ satisfies property $R(\epsilon s/2, M)$, then $\Pr_F[res2(F) < M/2] < 2p(F)$ by Theorem 14. Since $p(F)$ is $o(1)$ according to Lemma 32, the Theorem would follow.

It remains to prove that $\mathcal{B}_t$ satisfies property $R(\epsilon s/2, M)$. In the following, we think of ρ as drawn from $\mathcal{B}_t$. We let $\rho = ((x_1, v_1), \ldots, (x_r, v_r))$.

A 2-disjunction is *large* if it contains at least $d = \epsilon s/2$ literals, otherwise it is *small*. A 2-disjunction is *wide* if it contains at least $w = t/2(\log(t))^2$ free literals, otherwise it is *narrow*. We say that (x_i, v_i) *knocks* a 2-disjunction if it makes it true. We say that (x_i, v_i) *hits* a 2-disjunction if it makes true a literal in it. Notice that every knock is a hit, but a hit might not be a knock. We say that (x_i, v_i) is a *bad choice* if it does not knock the 2-disjunction but could have knocked it just by giving the opposite value to the variable. For $i \leq r$, we let ρ_i be $((x_1, v_1), \ldots, (x_i, v_i))$. When possible we simplify 2-disjunctions: we substitute subformulas of the form $l \vee (l \wedge l')$ by l in some order. We aim for a proof that

$$\Pr[C|_\rho \text{ is large}] \leq e^{-\frac{n^{1/3}}{(\log(n))^4}}, \tag{4.7}$$

where C is an arbitrary simplified 2-disjunction.

Let A_i be the event that $C|_{\rho_i}$ contains at least d distinct literals. Let A be the event $A_{|\rho|}$.

$$\Pr[A] \;=\; \Pr[A \wedge |\rho| < t/2] + \Pr[A \wedge |\rho| \geq t/2]. \tag{4.8}$$

Obviously $\Pr[A \wedge |\rho| < t/2] \leq \Pr[|\rho| < t/2]$ which is smaller than $e^{-t/8}$ by Chernoff bounds, so

$$(4.8) \leq e^{-\frac{n^{2/3}}{\log(n)}} + \Pr[A \mid |\rho| \geq t/2].$$

We show now that $\Pr[A \mid |\rho| \geq t/2]$ is expone ntially small. For every i such that $t/4 \leq i \leq t/2$, let B_i be the event that $C|_{\rho_i}$ is narrow, that is, it contains less than w free literals. Let D be the event that $|\rho| \geq t/2$. Then,

$$\Pr[A \mid D] \;=\; \Pr\left[A \wedge \bigvee_{j=t/4}^{t/2} B_j \;\middle|\; D\right] + \Pr\left[A \wedge \bigwedge_{j=t/4}^{t/2} \overline{B_j} \;\middle|\; D\right]. \tag{4.9}$$

We show that both terms in (4.9) are exponentially small. For every i such that $t/4 \leq i \leq t/2$, let T_i be the indicator random variable for the event that (x_i, v_i) is a knock. Then the second term in (4.9) is

$$\Pr\left[A \wedge \bigwedge_{j=t/4}^{t/2} \overline{B_j} \;\middle|\; D\right] \;\leq\; \Pr\left[\bigwedge_{i>t/4}^{t/2} T_i = 0 \wedge \bigwedge_{j=t/4}^{t/2} \overline{B_j} \;\middle|\; D\right] =$$

$$
\begin{aligned}
&= \prod_{i>t/4}^{t/2} \Pr\left[T_i = 0 \ \wedge \bigwedge_{j=t/4}^{t/2} \overline{B_j} \ \middle|\ \bigwedge_{j>t/4}^{i-1} T_j = 0 \wedge D\right] \leq \\
&= \prod_{i>t/4}^{t/2} \Pr\left[T_i = 0 \wedge \overline{B_{i-1}} \ \middle|\ \bigwedge_{j>t/4}^{i-1} T_j = 0 \wedge D\right] \leq \\
&= \prod_{i>t/4}^{t/2} \Pr\left[T_i = 0 \ \middle|\ \overline{B_{i-1}} \wedge \bigwedge_{j>t/4}^{i-1} T_j = 0 \wedge D\right] \leq \\
&\leq \prod_{i>t/4}^{t/2} \left(1 - \frac{w}{2(n-i+1)}\right) \leq \left(1 - \frac{w}{2n}\right)^{t/4} \leq \\
&\leq e^{-\frac{wt}{8n}} = e^{-\frac{n^{1/3}}{(\log(n))^3}}.
\end{aligned}
$$

The first term in (4.9) is also exponentially small. Observe that

$$
\Pr\left[A \wedge \bigvee_{j=t/4}^{t/2} B_j \ \middle|\ D\right] = \Pr\left[\bigvee_{j=t/4}^{t/2} (A \wedge B_j) \ \middle|\ D\right] = \tag{4.10}
$$

$$
\leq \sum_{j=t/4}^{t/2} \Pr\left[A_j \wedge B_j \mid D\right]. \tag{4.11}
$$

The last inequality is true because A implies A_j for any $j \leq t/2$.

Lemma 33 *If j is such that $t/4 \leq j \leq t/2$, then* $\Pr\left[A_j \wedge B_j \mid D\right] \leq e^{-\frac{n^{2/3}}{(\log(n))^6}}$.

Proof: For every $i \leq j$ let S_i be the indicator random variable for the event that (x_i, v_i) hits $C|_{\rho_{i-1}}$, that is, that (x_i, v_i) gives value $true$ to a literal in $C|_{\rho_{i-1}}$. Let $S = \sum_{i=1}^{j} S_i$. We divide the calculation in two parts: what happens when the number of hits is less than a certain $h = t/(\log(t))^2$ and what happens otherwise.

$$
\Pr\left[A_j \wedge B_j \mid D\right] = \Pr\left[A_j \wedge B_j \wedge S < h \mid D\right] + \Pr\left[A_j \wedge B_j \wedge S \geq h \mid D\right]
$$

We start by the easiest part. The intuition is that if the 2-disjunction is large it would be extremely difficult to hit it only a few times.

Sublemma 1 $\Pr\left[A_j \wedge S < h \mid D\right] \leq e^{-\frac{n^{2/3}}{(\log(n))^2}}$.

Proof: Let $Y = \{(a_1, \ldots, a_j) \in \{0,1\}^j : \sum_{i=1}^{j} a_i < h\}$. Observe that A_j implies A_i for every $i \leq j$ because if $C|_{\rho_j}$ is large, so is $C|_{\rho_i}$. Then,

$$
\Pr\left[A_j \wedge S < h \mid D\right] = \Pr\left[A_j \wedge \sum_{i=1}^{j} S_i < h \ \middle|\ D\right] =
$$

$$
\begin{aligned}
&= \sum_{\overline{a}\in Y} \Pr\left[A_j \wedge \bigwedge_{i=1}^{j} S_i = a_i \;\middle|\; D\right] = \\
&= \sum_{\overline{a}\in Y} \prod_{i=1}^{j} \Pr\left[A_j \wedge S_i = a_i \;\middle|\; \bigwedge_{k=1}^{i-1} S_k = a_k \wedge D\right] \leq \\
&\leq \sum_{\overline{a}\in Y} \prod_{i=1}^{j} \Pr\left[A_{i-1} \wedge S_i = a_i \;\middle|\; \bigwedge_{k=1}^{i-1} S_k = a_k \wedge D\right] \leq \\
&\leq \sum_{\overline{a}\in Y} \prod_{i=1}^{j} \Pr\left[S_i = a_i \;\middle|\; A_{i-1} \wedge \bigwedge_{k=1}^{i-1} S_k = a_k \wedge D\right].
\end{aligned}
$$

Fix $i \in \{1, \ldots, j\}$.

$$
\Pr\left[S_i = 1 \;\middle|\; A_{i-1} \wedge \bigwedge_{k=1}^{i-1} S_k = a_k \wedge D\right] \geq \frac{d}{2(n-i+1)} \geq \frac{d}{2n}.
$$

Since there are at least $j-h$ zeros in $(a_1, \ldots, a_j)$, we obtain

$$
\begin{aligned}
\Pr\left[A_j \wedge S < h \mid D\right] &\leq \sum_{\overline{a}\in Y} \left(1 - \frac{d}{2n}\right)^{j-h} \leq \\
&\leq \sum_{i<h} \binom{j}{i} e^{-\frac{d(j-h)}{2n}} \leq hj^h e^{-\frac{d(j-h)}{2n}} \leq \\
&\leq \exp\left(-\frac{d(j-h)}{2n} + h\log(j) + \log(h)\right) \leq \\
&\leq e^{-\frac{dt}{10n} + \frac{t}{\log(t)}} \leq e^{-\frac{n^{2/3}}{(\log(n))^2}}.
\end{aligned}
$$

□ (of sublemma 1)

The last thing to do is to see what happens when the number of hits is big.

Sublemma 2 $\Pr\left[A_j \wedge B_j \wedge S \geq h \mid D\right] \leq e^{-\frac{n^{2/3}}{(\log(n))^5}}$.

Proof: For every $1 \leq i \leq t/2$, let $T_i \in \{\mathrm{k}, \mathrm{b}, \mathrm{n}\}$ be a random variable indicating whether (x_i, v_i) is a knock, a bad choice, or none of the previous respectively for $C|_{\rho_{i-1}}$. For $t \in \{\mathrm{k}, \mathrm{b}, \mathrm{n}\}$, let S_i^t be the indicator random variable for the event that $T_i = t$, and let $S^t = \sum_{i=1}^{j} S_i^t$. Thus, S^{k} is the number of knocks and S^{b} is the number of bad choices of ρ_j. For the rest of the proof we will skip the condition on D and the subindices from A and B. Fix ρ satisfying $A \wedge B \wedge S \geq h$. Note that the number of knocks is 0 because the 2-disjunction still exists, so $S^{\mathrm{k}} = 0$. Now let be $b = (h-w)/2$, we now claim that $S^{\mathrm{b}} \geq b$. Suppose for contradiction that the number of bad choices is less than b. Every bad choice (x_i, v_i) removes at most one free literal. Moreover, since there are no knocks, every hit (x_i, v_i) that is not a bad choice increases the number of free literals by at least one. The reason is that such a hit turns a conjunction into a free literal. Remember that we simplify the 2-disjunction when possible, and so the literal was not free before the hit (x_i, v_i) is applied. It follows then that the number of free literals in $C|_{\rho_j}$ is at least

$$(S - S^{\mathrm{b}}) - S^{\mathrm{b}} > h - 2b = w,$$

a contradiction with the fact that B holds under ρ.

So far we have proved that

$$\Pr[A \wedge B \wedge S \geq h] \leq \Pr[S^{\mathrm{k}} = 0 \wedge S^{\mathrm{b}} \geq b].$$

The intuition behind why this last probability is small is that every bad choice could have been a knock. This makes it unlikely that ρ produces many bad choices and no knocks. In what follows, we will prove this intuition using martingales.

Claim 3 $\Pr[S^{\mathrm{k}} = 0 \wedge S^{\mathrm{b}} \geq b] \leq e^{-\frac{n^{2/3}}{(\log(n))^5}}$.

Proof: For $t \in \{\mathrm{k}, \mathrm{b}, \mathrm{n}\}$ and $i \in \{1, \ldots, j\}$, let $P_i^t = \Pr[T_i = t \mid \rho_0, \ldots, \rho_{i-1}]$. Note that P_i^t is a random variable. We define a martingale $X_0, \ldots, X_j$ with respect to $\rho_0, \ldots, \rho_j$ as follows: Let

$$\begin{aligned} X_0 &= 0, \\ X_{i+1} &= X_i + S_{i+1}^{\mathrm{b}} - P_{i+1}^{\mathrm{b}}. \end{aligned}$$

Recall that S_{i+1}^{b} is the indicator random variable for the event that $T_{i+1} = \mathrm{b}$. Observe that

$$\begin{aligned} \mathrm{E}[X_{i+1} \mid \rho_0, \ldots, \rho_i] &= (X_i + 1 - P_{i+1}^{\mathrm{b}}) \cdot P_{i+1}^{\mathrm{b}} + (X_i - P_{i+1}^{\mathrm{b}}) \cdot (1 - P_{i+1}^{\mathrm{b}}) = \\ &= (X_i - P_{i+1}^{\mathrm{b}})(P_{i+1}^{\mathrm{b}} + 1 - P_{i+1}^{\mathrm{b}}) + P_{i+1}^{\mathrm{b}} = X_i. \end{aligned}$$

Hence, $\{X_i\}_i$ is a martingale with respect to $\{\rho_i\}_i$. Observe also that $X_j = S^{\mathrm{b}} - \sum_{i=1}^{j} P_i^{\mathrm{b}}$. Similarly, we define $Y_0, \ldots, Y_j$ as follows: Let $Y_0 = 0$, and $Y_{i+1} = Y_i + S_{i+1}^{\mathrm{k}} - P_{i+1}^{\mathrm{k}}$. It is also easy to see that $\{Y_i\}_i$ is a martingale with respect to $\{\rho_i\}_i$. Again, $Y_j = S^{\mathrm{k}} - \sum_{i=1}^{j} P_i^{\mathrm{k}}$.

We will use the following form of Azuma's Inequality: Let $X_0, \ldots, X_n$ be a martingale such that $|X_i - X_{i-1}| \leq 1$; then, $\Pr[|X_n - X_0| \geq \lambda] \leq 2e^{-\lambda^2/n}$ for every $\lambda > 0$. In the next calculation we will also use the fact that $P_i^{\mathrm{k}}(\rho) = P_i^{\mathrm{b}}(\rho)$ for every ρ and $i \in \{1, \ldots, j\}$.

$$\begin{aligned} \Pr[S^{\mathrm{k}} = 0 \wedge S^{\mathrm{b}} \geq b] &= \Pr[S^{\mathrm{k}} = 0 \wedge S^{\mathrm{b}} \geq b \wedge X_j \geq b/2] + \\ &+ \Pr[S^{\mathrm{k}} = 0 \wedge S^{\mathrm{b}} \geq b \wedge X_j < b/2]. \end{aligned}$$

The first summand is bounded by $\Pr[X_j \geq b/2] \leq 2e^{-b^2/4j}$ by Azuma's Inequality. The second summand is bounded by

$$\begin{aligned} \Pr\left[S^{\mathrm{k}} = 0 \wedge \textstyle\sum_{i=1}^{j} P_i^{\mathrm{b}} \geq b/2\right] &\leq \Pr\left[S^{\mathrm{k}} = 0 \wedge \textstyle\sum_{i=1}^{j} P_i^{\mathrm{k}} \geq b/2\right] \leq \\ &\leq \Pr[Y_j \leq -b/2] \leq \\ &\leq 2e^{-b^2/4j}, \end{aligned}$$

by Azuma's Inequality again. Therefore, the sum is bounded by

$$4e^{-\frac{t}{32(\log(t))^4}} \leq e^{-\frac{n^{2/3}}{(\log(n))^5}}$$

as required. $\square$ (of claim 3 and sublemma 2).

With both sublemmas proved, so is Lemma 33. We are ready to complete the proof of our goal (4.7). We have shown that

$$\Pr\left[C|_\rho \text{ is large}\right] \;\leq\; e^{-\frac{n^{2/3}}{\log(n)}} + e^{-\frac{n^{1/3}}{(\log(n))^3}} + \sum_{i=t/4}^{t/2} e^{-\frac{n^{2/3}}{(\log(n))^6}} \leq e^{-\frac{n^{1/3}}{(\log(n))^4}}.$$

$\square$

4.4 Relationship to the Calculus of Lovász-Schrijver

It has been long known that Boolean clauses may be rewritten as linear inequalities over $\{0,1\}$. Let C be a clause of the form

$$x_{i_1} \vee \ldots \vee x_{i_r} \vee \neg x_{j_1} \vee \ldots \vee \neg x_{j_s}$$

The translation of C, denoted by $t(C)$, is a linear polynomial on $x_1, \ldots, x_n$ defined as follows:

$$x_{i_1} + \cdots + x_{i_r} + (1 - x_{j_1}) + \cdots + (1 - x_{j_s}).$$

Clearly, a truth assignment to the variables $x_1, \ldots, x_n$ satisfies C if and only if $t(C) \geq 1$. As usual, *true* and *false* are represented by 1 and 0 respectively.

Just as clauses correspond to linear inequalities, 2-disjunctions correspond to quadratic inequalities through the following straightforward translation. Let D be a 2-disjunction of the form

$$(l_{1,1} \wedge l_{1,2}) \vee \ldots \vee (l_{s,1} \wedge l_{s,2})$$

where each $l_{i,j}$ is a literal in the variables $x_1, \ldots, x_n$. The translation of D, denoted by $t(D)$, is a degree-two polynomial on $x_1, \ldots, x_n$ defined as follows:

$$t(l_{1,1})t(l_{1,2}) + \cdots + t(l_{s,1})t(l_{s,2})$$

where $t(l) = x_i$ if $l = x_i$ and $t(l) = 1 - x_i$ if $l = \neg x_i$. Again, a truth assignment satisfies D if and only if $t(D) \geq 1$.

We turn next to the definition of the Cutting Planes proof system, denoted by CP. The allowed formulas are linear inequalities with integer coefficients and there only are two rules of inference:

Addition:
$$\frac{a_1x_1 + \cdots + a_nx_n \geq a_0 \qquad b_1x_1 + \cdots + b_nx_n \geq b_0}{(a_1+b_1)x_1 + \cdots + (a_n+b_n)x_n \geq a_0 + b_0}$$

Integer Division:
$$\frac{(b \cdot a_1)x_1 + \cdots + (b \cdot a_n)x_n \geq a_0}{a_1x_1 + \cdots + a_nx_n \geq \lceil a_0/b \rceil}$$

Here $a_0, a_1, \ldots, a_n, b_0, b_1, \ldots, b_n$ and b are all integers. We also allow the axioms $x_i \geq 0$ and $1 - x_i \geq 0$. The goal of the system is to refute a given set of linear inequalities by deriving $0 \geq 1$. It is well-known that CP polynomially simulates Resolution when clauses are presented as linear inequalities according to the translation above [CCT87].

We present an extension of CP to work with quadratic inequalities. Pudlák called it CP^2 in [Pud01]. In this system, the allowed formulas are the quadratic inequalities with integer coefficients. The only rules of inference are addition and integer division for quadratic inequalities, in addition to the axioms

$$\begin{array}{ll} x_i \geq 0 & 1 - x_i \geq 0 \\ x_ix_j \geq 0 & x_i - x_ix_j \geq 0 \\ x_j - x_ix_j \geq 0 & 1 - x_i - x_j + x_ix_j \geq 0. \end{array}$$

The first goal of this section is to show that CP^2 polynomially simulates Res(2). This simple observation was made in joint work of the author with Bonet and Levy [ABL01]. We start with a simple lemma.

Lemma 34 *For every* 2*-disjunction* D*, the inequality* $t(D) \geq 0$ *has a CP*^{2}*-proof of size* $O(|D|)$*.*

Proof: For each pair of literals l_1 and l_2, the inequality $t(l_1)t(l_2) \geq 0$ is exactly one of the axioms of CP^2. The rule of addition gives the desired proof of $t(D) \geq 0$. □

Theorem 16 *Let* $\{D_1, \ldots, D_s\}$ *be a set of* 2*-disjunctions with a Res(2)-refutation of size* S*. Then, the system of inequalities* $\{t(D_1) \geq 1, \ldots, t(D_s) \geq 1\}$ *has a CP*^{2}*-refutation of size* $O(S^2)$*.*

Proof: Take the Res(2)-refutation of size S and replace each 2-disjunction D by the inequality $t(D) \geq 1$. We show how to simulate each of the rules of Res(2).

Weakening: Suppose that $D \vee E$ is derived by the rule of weakening from D, where E is an arbitrary 2-disjunction. We get $t(D) + t(E) \geq 1$ from $t(D) \geq 1$ by deriving $t(E) \geq 0$ separately and adding.

Introduction of conjunction: Suppose that $D \vee (l_1 \wedge l_2)$ is derived by the rule of introduction of conjunction from $D_1 \vee l_1$ and $D_2 \vee l_2$. We get $t(D) + t(l_1)t(l_2) \geq 1$ from $t(D_1) + t(l_1) \geq 1$

and $t(D_2) + t(l_2) \geq 1$ as follows. Add the two inequalities with the rule of addition. This gives $t(D_1) + t(D_2) + t(l_1) + t(l_2) \geq 2$. Derive the inequality $(1 - t(l_1))(1 - t(l_2)) \geq 0$ separately and add to obtain $t(D_1)+t(D_2)+t(l_1)+t(l_2)+(1-t(l_1))(1-t(l_2)) \geq 2$. Expanding and rearranging, we obtain $t(D_1)+t(D_2)+t(l_1)t(l_2) \geq 1$. It remains to see how to eliminate duplicate terms to get $t(D) + t(l_1)t(l_2) \geq 1$. For this, it suffices to add the easily derived inequality $t(r_1)t(r_2) \geq 0$ for each term $t(r_1)t(r_2)$ in $t(D)+t(l_1)t(l_2)$ that is not already duplicated in $t(D_1)+t(D_2)+t(l_1)t(l_2)$. Rearranging, this gives $2(t(D) + t(l_1)t(l_2)) \geq 1$. A single application of the rule of division gives the desired inequality.

Cut: Suppose that D is derived by a cut from $D_1 \vee (l_1 \wedge l_2)$ and $D_2 \vee \neg l_1 \vee \neg l_2$. We get $t(D) \geq 1$ from $t(D_1) + t(l_1)t(l_2) \geq 1$ and $t(D_2) + 1 - t(l_1) + 1 - t(l_2) \geq 1$ as follows. Derive $t(D_2) \geq 0$ separately and add it to one copy of the second inequality and two copies of the first. This gives $2t(D_1) + 2t(D_2) + 2t(l_1)t(l_2) + 1 - t(l_1) + 1 - t(l_2) \geq 3$. Derive the inequality $t(l_1)(1 - t(l_2)) + t(l_2)(1 - t(l_1)) \geq 0$ separately and add to obtain

$$2t(D_1) + 2t(D_2) + 2t(l_1)t(l_2) + 1 - t(l_1) + 1 - t(l_2) + t(l_1)(1 - t(l_2)) + t(l_2)(1 - t(l_1)) \geq 3.$$

Expanding, we observe that this is the same as $2t(D_1) + 2t(D_2) \geq 1$. The rule of division gives then $t(D_1) + t(D_2) \geq 1$. Finally, we apply the same technique as before to eliminate duplicates from $t(D_1) + t(D_2)$ and obtain $t(D) \geq 1$ as required.

It remains to see how to simulate the axioms of Res(2) of the type $x \vee \neg x$. These get translated into $x + 1 - x \geq 1$ and we do not know how to derive that. However, we may proceed in a different way. Consider the first time that a non-weakening rule is applied over a line that is derived from an axiom by weakening. Suppose that the rule is a cut, say over $D \vee (l_1 \wedge l_2)$ and $A \vee \neg l_1 \vee \neg l_2$, where $A \vee \neg l_1 \vee \neg l_2$ is derived from $x \vee \neg x$ by weakening. If both $\neg l_1$ and $\neg l_2$ were introduced by weakening, then the result of the cut will contain $x \vee \neg x$ as a subformula. The translation will be of the form of Lemma 34 and will be easy to derive. If only one of $\neg l_1$ and $\neg l_2$ is introduced by weakening, say $\neg l_2$, and the other is one of x or $\neg x$, say $\neg x$, then the cut is simulated by an addition with $x(1 - l_2) \geq 0$ which is an axiom, and some additional weakening in the form of Lemma 34. Suppose next that the rule is a cut, say over $D \vee \neg l_1 \vee \neg l_2$ and $A \vee (l_1 \wedge l_2)$ where $A \vee (l_1 \wedge l_2)$ is derived from $x \vee \neg x$ by weakening. Obviously, $l_1 \wedge l_2$ is introduced by weakening and so the result of the cut contains $x \vee \neg x$. There are two cases left: an application of the rule of introduction of conjunction with a principal formula introduced by weakening over $x \vee \neg x$, and introduction of conjunction with one of x or $\neg x$ as a principal formula, say x. The former is as before: the result of the rule contains $x \vee \neg x$. The latter is also easy: we simulate the introduction of the conjunction $x \wedge l_2$ by an addition with $xl_2 + 1 - x - l_2 \geq 0$ which is an axiom whatever l_2 is, and additional weakenings in the form of Lemma 34. This completes the proof of the Theorem. $\square$

The calculus of Lovász-Schrijver, denoted by LS, was also introduced by Pudlák [Pud97] and gave it this name to credit the seminal paper [LS91]. In this system, the allowed formulas are also the quadratic inequalities but the rules are different. Namely, the only rules permitted are addition of arbitrary inequalities and multiplication of linear inequalities by $x_i \geq 0$ or $1 - x_i \geq 0$:

Multiplication:

$$\frac{L \geq 0}{Lx_i \geq 0} \qquad\qquad \frac{L \geq 0}{L(1 - x_i) \geq 0}$$

where L is a linear polynomial on the variables $x_1, \ldots, x_n$. Again, we allow the axioms $x_i \geq 0$ and $1 - x_i \geq 0$, and the axiom $x_i^2 - x_i \geq 0$ to enforce 0/1-solutions. It is not known whether LS polynomially simulates Res(2). In fact, a recent result of the author in joint work with Bonet [AB01] implies that the positive answer would have important consequences for the automatization of Resolution. However, if we allow the rule of division for quadratic inequalities, the resulting system does polynomially simulate Res(2).

Theorem 17 *Let $\{D_1, \ldots, D_s\}$ be a set of 2-disjunctions with a Res(2)-refutation of size S. Then, the system of inequalities $\{t(D_1) \geq 1, \ldots, t(D_s) \geq 1\}$ has a refutation in LS extended by the rule of division for quadratic inequalities of size $O(S)$.*

Proof: The axioms of CP^2 are easily derivable with the rule of Mutiplication applied on the axioms $x_i \geq 0$ and $1 - x_i \geq 0$. The result follows by Theorem 16. □

We present yet another proof system that polynomially simulates Res(2). The goal is to avoid the rule of division at all. The new system is called Q. The allowed formulas are polynomial inequalities of degree at most four. The rules of inference are addition of arbitrary inequalities and multiplication of quadratic inequalities:

Addition:

$$\frac{P_1 \geq 0 \qquad P_2 \geq 0}{P_1 + P_2 \geq 0}$$

Multiplication:

$$\frac{Q_1 \geq 0 \qquad Q_2 \geq 0}{Q_1 Q_2 \geq 0}$$

where P_1 and P_2 are arbitrary polynomials, and Q_1 and Q_2 are quadratic polynomials. We also allow the axioms $x_i \geq 0$ and $1 - x_i \geq 0$, and $x_i^2 - x_i \geq 0$ to enforce 0/1-solutions.

Lemma 35 *For every 2-disjunction D, the inequality $t(D) \geq 0$ has a Q-proof of size $O(|D|)$.*

Proof: This is straightforward as in Lemma 34. Indeed, $t(l_1)t(l_2) \geq 0$ is simply the result of multiplying $t(l_1) \geq 0$ and $t(l_2) \geq 0$, which are axioms by definition. The rule of addition gives the result. □

Lemma 36 *Let A be a quadratic polynomial, and let l_1 and l_2 be literals. The inequality $A + t(l_1)t(l_2) \geq 1$ has a Q-proof of size $O(|A|)$ from the inequalities $A + 2t(l_1)t(l_2) \geq 1$ and $A \geq 0$.*

Proof: In the following, let $z_1 = t(l_1)$ and $z_2 = t(l_2)$. Consider the following sequence of inequalities: Start with $A + 2z_1z_2 \geq 1$ and rearrange to obtain

$$A + 2z_1z_2 - 1 \geq 0. \tag{4.12}$$

Derive $(1 - z_1)z_2 \geq 0$ and multiply by (4.12) to obtain

$$(A + 2z_1z_2 - 1)(1 - z_1)z_2 \geq 0. \tag{4.13}$$

Multiply (4.12) by $(1 - z_2) \geq 0$ to obtain

$$(A + 2z_1z_2 - 1)(1 - z_2) \geq 0. \tag{4.14}$$

Add (4.13) and (4.14) to obtain

$$A + 3z_1z_2 - 1 - Az_1z_2 - 2z_1^2z_2^2 \geq 0 \tag{4.15}$$

after rearranging. Then derive $z_1z_2 \geq 0$ and multiply with $A \geq 0$ (which is given) to obtain $Az_1z_2 \geq 0$. Add this to (4.15) to obtain

$$A + 3z_1z_2 - 2z_1^2z_2^2 - 1 \geq 0. \tag{4.16}$$

Derive $z_2z_2 \geq 0$ and multiply by the axiom $z_1^2 - z_1 \geq 0$ to obtain $(z_1^2 - z_1)z_2z_2 \geq 0$. Add two copies of this to (4.16) and rearrange to obtain

$$A + 3z_1z_2 - 2z_1z_2^2 - 1 \geq 0. \tag{4.17}$$

Similarly, multiply the axioms $z_2^2 - z_2 \geq 0$ and $z_1 \geq 0$ to obtain $(z_2^2 - z_2)z_1 \geq 0$. Add two copies of this to (4.17) and rearrange to obtain $A + z_1z_2 \geq 1$ as required. □

We conclude with the promised simulation of Res(2) by Q.

Theorem 18 *Let $\{D_1, \ldots, D_s\}$ be a set of 2-disjunctions with a Res(2)-refutation of size S. Then, the system of inequalities $\{t(D_1) \geq 1, \ldots, t(D_s) \geq 1\}$ has a Q-refutation of size $O(S^2)$.*

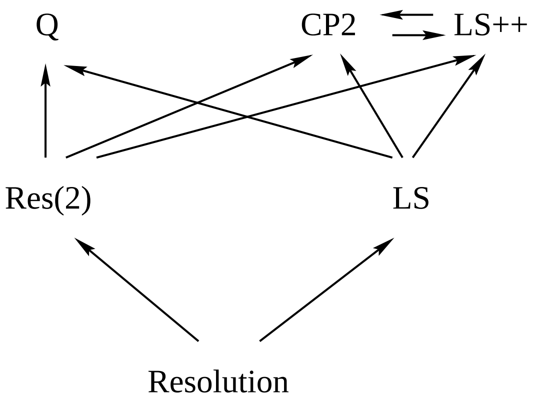

Figure 4.1: Simulation results between proof systems.

Proof: We first observe that the axioms of CP^2 are easily derivable in Q. Now, inspection on the proof of Theorem 16 reveals that the rule of division is only used in the following form: from $A + 2t(l_1)t(l_2) \geq 1$, derive $A + t(l_1)t(l_2) \geq 1$ with A a quadratic polynomial. This can be done efficiently in Q by Lemma 36 if $A \geq 0$ is given. However, the A's in the proof of Theorem 16 are always additions of terms of the form $t(l_1)t(l_2)$, and so $A \geq 0$ is easily derivable by Lemma 35. □

Diagram 4.1 is a summary of the known polynomial simulation results between the proof systems considered in this section. Recall that LS denotes the Lovász-Schrijver proof system. We will denote by LS++ the extension of Lovász-Schrijver by the rule of division for quadratic inequalities. An arrow between system P_1 and system P_2 means that P_2 polynomially simulates P_1.

Chapter 5

Improved Bounds on the WPHP and Infinitely Many Primes

5.1 The Proof in Propositional Logic

The propositional form of the Weak Pigeonhole Principle PHP^m_n that we use is formalized by the following sequent:

$$\bigwedge_{i=1}^{m} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i \neq j}}^{m} p_{i,k} \wedge p_{j,k}.$$

Our goal is to prove a size-depth trade-off upper bound for proofs of PHP^{2n}_n. As mentioned in the introduction, our technique consists in reducing PHP^{2n}_n to $\mathrm{PHP}^{n^2}_n$. The difference with previous reductions is that our composition of mappings is done efficiently mimicking the proofs of Savitch's and Nepomnjaščij's Theorems in Complexity Theory. To complete the proof, we will use the fact that $\mathrm{PHP}^{n^2}_n$ has LK proofs of size $n^{O(\log^{(d)}(n))}$ and depth $O(d)$, where $\log^{(d)}(n)$ is the d-wise composition of $\log$ with itself [PWW88, Kra95].

For the sake of clarity of exposition, it is more convenient to prove the following extreme case of the trade-off first.

Theorem 19 PHP^{2n}_n *has* LK *proofs of size* $n^{O(\log\log(n))}$ *and depth* $3\log\log(n)$.

Proof: Let $d = \log\log(n)$. For every $\alpha \in \{0,1\}^{\leq d}$, we define numbers L_α and R_α inductively as follows: Let $L_\lambda = \log(n)$, $R_\lambda = 0$, and $L_{\alpha 0} = L_\alpha$, $R_{\alpha 1} = R_\alpha$,

$$L_{\alpha 1} = R_{\alpha 0} = \frac{1}{2}(L_\alpha + R_\alpha).$$

Now we define sets A_α, B_α and C_α as follows: Let $A_\alpha = \{0,\ldots,2^{L_\alpha}n-1\}$, $B_\alpha = \{0,\ldots,2^{R_\alpha}n-1\}$ and $C_\alpha = \{0,\ldots,2^{\frac{1}{2}(L_\alpha+R_\alpha)}n-1\}$. Observe that $A_\lambda = \{0,\ldots,n^2-1\}$ and $B_\lambda = \{0,\ldots,n-1\}$.

For $\alpha \in \{0,1\}^d$, $x \in A_\alpha$, $y \in B_\alpha$, let R^α_{xy} be defined as follows: If $\lfloor x/2n \rfloor \neq \lfloor y/n \rfloor$, define $R^\alpha_{xy} = 0$; otherwise, define $R^\alpha_{xy} = P_{x'y'}$ where $x' = x \mod 2n$ and $y' = y \mod n$. For $\alpha \in \{0,1\}^{<d}$, $x \in A_\alpha$, $y \in B_\alpha$, define

$$R^\alpha_{xy} = \bigvee_{z \in C_\alpha} (R^{\alpha 0}_{xz} \wedge R^{\alpha 1}_{zy}).$$

It is easy to see that the size of R^α_{xy} is bounded by $n^{2(d-|\alpha|)}$ and the depth is bounded by $2(d-|\alpha|)$. In particular, R^λ_{xy} has size bounded by $n^{2\log\log(n)}$ and depth bounded by $2\log\log(n)$.

In the following, we will abbreviate some expressions by ignoring the conjunction symbol. Thus, if F and G are Boolean formulas, the notation FG will be used instead of $F \wedge G$, and similarly for longer conjunctions. This will allow us ignore some parenthesis following the rules of arithmetic when $\vee$ and $\wedge$ are interpreted as $+$ and $\times$ respectively.

We want to prove the following sequents

$$\bigwedge_{\alpha \in \{0,1\}^k} \bigwedge_x \bigvee_y R^\alpha_{xy} \rightarrow \bigvee_{\alpha \in \{0,1\}^k} \bigvee_y \bigvee_{x_1 \neq x_2} R^\alpha_{x_1 y} R^\alpha_{x_2 y} \tag{5.1}$$

in size bounded by n^{ck} and depth bounded by $3k$, where c is a sufficiently large constant independent of n ($c = 20$ should work). When $k = d = \log\log(n)$, it is easy to see that sequent (5.1) is equivalent to PHP^{2n}_n after contraction of repeated formulas, and the theorem will follow. We fix a sufficiently large n, and proceed by induction on k.

Observe that the base case $k = 0$ is precisely the sequent $\mathrm{PHP}^{n^2}_n(R^\lambda)$. Since the sequent $\mathrm{PHP}^{n^2}_n$ has LK proofs of size $n^{O(\log^{(3)}(n))}$ and constant depth, it follows that $\mathrm{PHP}^{n^2}_n(R^\lambda)$ has LK-proofs of size bounded by $n^{c\log\log(n)}$ and depth bounded by $3\log\log(n)$. Here is where we need n to be sufficiently large.

Suppose next that we have proved sequent (5.1) for a $k > 0$ in size n^{ck} and depth $3k$. We prove it for $k+1$. We first prove the following sequents for every $\alpha \in \{0,1\}^k$ and w:

$$\bigwedge_x \bigvee_y R^{\alpha 0}_{xy} \wedge \bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \rightarrow \bigvee_y R^\alpha_{wy}. \tag{5.2}$$

Recall that R^α_{wy} stands for $\bigvee_z R^{\alpha 0}_{wz} R^{\alpha 1}_{zy}$. Start with the sequents $R^{\alpha 0}_{wz} \rightarrow R^{\alpha 0}_{wz}$ and $\bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \rightarrow \bigvee_y R^{\alpha 1}_{zy}$, and apply right $\wedge$ introduction to obtain

$$R^{\alpha 0}_{wz}, \bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \rightarrow R^{\alpha 0}_{wz} \wedge \bigvee_y R^{\alpha 1}_{zy}. \tag{5.3}$$

By distributivity we easily get

$$R^{\alpha 0}_{wz}, \bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \rightarrow \bigvee_y R^{\alpha 0}_{wz} R^{\alpha 1}_{zy}. \tag{5.4}$$

By left $\vee$-introduction, left weakening, left $\wedge$-introduction, right $\vee$-introduction and commutativity of $\vee$, in this order, we get the desired sequent (5.2).

Next we prove the following sequents for every $\alpha \in \{0,1\}^k$, $w_1 \neq w_2$ and y:

$$R^{\alpha}_{w_1 y} R^{\alpha}_{w_2 y} \to \bigvee_z \bigvee_{x_1 \neq x_2} R^{\alpha 0}_{x_1 z} R^{\alpha 0}_{x_2 z} \vee \bigvee_z \bigvee_{x_1 \neq x_2} R^{\alpha 1}_{x_1 z} R^{\alpha 1}_{x_2 z}. \tag{5.5}$$

Recall that $R^{\alpha}_{w_i y}$ stands for $\bigvee_z R^{\alpha 0}_{w_i z} R^{\alpha 1}_{zy}$. Using distributivity, derive the sequent

$$R^{\alpha}_{w_1 y} R^{\alpha}_{w_2 y} \to \bigvee_{z_1, z_2} R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y}. \tag{5.6}$$

For $z_1 = z_2$, derive the sequent $R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y} \to R^{\alpha 0}_{w_1 z_1} R^{\alpha 0}_{w_2 z_2}$. For $z_1 \neq z_2$, derive the sequent $R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y} \to R^{\alpha 1}_{z_1 y} R^{\alpha 1}_{z_2 y}$. Left $\vee$-introduction and right $\vee$-introduction gives the sequent

$$\bigvee_{z_1, z_2} R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y} \to \bigvee_{z_1 = z_2} R^{\alpha 0}_{w_1 z_1} R^{\alpha 0}_{w_2 z_2} \vee \bigvee_{z_1 \neq z_2} R^{\alpha 1}_{z_1 y} R^{\alpha 1}_{z_2 y}. \tag{5.7}$$

A cut with sequent (5.6), right weakening, right $\vee$-introduction and commutativity of $\vee$, in this order, give the desired sequent (5.5).

Now combine all sequents (5.2) by right $\wedge$-introduction, left $\wedge$-introduction and commutativity of $\wedge$, in this order, to obtain

$$\bigwedge_{\alpha \in \{0,1\}^{k+1}} \bigwedge_x \bigvee_y R^{\alpha}_{xy} \to \bigwedge_{\alpha \in \{0,1\}^k} \bigwedge_x \bigvee_y R^{\alpha}_{xy}. \tag{5.8}$$

Similarly, combine all sequents (5.5) by left $\vee$-introduction, right $\vee$-introduction and commutativity of $\vee$, in this order, to obtain

$$\bigvee_{\alpha \in \{0,1\}^k} \bigvee_y \bigvee_{x_1 \neq x_2} R^{\alpha}_{x_1 y} R^{\alpha}_{x_2 y} \to \bigvee_{\alpha \in \{0,1\}^{k+1}} \bigvee_y \bigvee_{w_1 \neq w_2} R^{\alpha}_{w_1 y} R^{\alpha}_{w_2 y}. \tag{5.9}$$

Finally, two cuts using sequents (5.1), (5.8) and (5.9) give the desired result for $k+1$. If c is sufficiently large, it is easy to check that the size of this proof is bounded by $n^{c(k+1)}$ and its depth is bounded by $3(k+1)$. This completes the induction step. □

An analogous argument mimicking the proof of Nepomnjaščij's Theorem, instead of Savitch's Theorem as above, would give a general size-depth trade-off upper bound. Since the notation in the proof would get fairly tedious, we prefer to state it without proof and get it as a corollary to Theorem 24 below (see the end of Section 5.3).

Theorem 20 PHP^{2n}_n *has* LK *proofs of size* $n^{O(d(\log(n))^{2/d})}$ *and depth* d.

5.2 Nepomnjaščii's Theorem, Formalized

Let us briefly recall the proof of a general form of Nepomnjaščij's Theorem. This will be of help later.

Theorem 21 (General Form of Nepomnjaščij's Theorem) *Let $K = K(n)$, $T = T(n)$ and $S = S(n)$ be time-constructible functions such that $K(n) \geq 2$. For every non-deterministic Turing machine running in simultaneous time T and space S, there exists an equivalent alternating Turing machine running in time $O(SK\log(T)/\log(K))$ and $2\log(T)/\log(K)$ alternations.*

Proof: Let M be a non-deterministic Turing machine running in simultaneous time T and space S. The idea is to divide the reachability problem between configurations of M into many equivalent subproblems of smaller size. Hence, configuration C_K is reachable from configuration C_0 in N steps if and only if there exist $K-1$ intermediate configurations $C_1, \ldots, C_{K-1}$ such that for every $i \in \{0, \ldots, K-1\}$, configuration C_{i+1} is reachable from configuration C_i in N/K steps. An alternating Turing machine can existentially quantify those intermediate configurations, and universally branch to check that every two consecutive configurations are reachable from each other in the appropriate number of steps. Applying this recursively yields an alternating machine that checks whether an accepting configuration is reachable from the initial configuration of M.

The details of the calculations follow. Since M runs in space S, each configuration can be coded by a binary word of length $O(S)$. The outermost reachability problem has length T since M runs in time T. The second level of reachability subproblems have length T/K. In general, the subproblems at level i have length T/K^i. After $\log(T)/\log(K)$ levels, we reached a trivial reachability subproblem. Therefore, the whole computation of the simulating machine takes time $O(SK\log(T)/\log(K))$ and $2\log(T)/\log(K)$ alternations. $\square$

Our next goal is to formalize the ideas of Theorem 21 into a theorem of the bounded arithmetic theory $I\Delta_0$. We will use the beautiful arithmetization of sequences in $I\Delta_0$ of Chapter V, Section 3, in the book of Hájek and Pudlák [HP93]. In summary, the arithmetization allows us to manipulate sequences provably in $I\Delta_0$. Thus, there are formulas $Seq(s)$ meaning that s is the code of a sequence, $(s)_i = x$ meaning that the i-th element of the sequence s is x, and $s \frown t = p$ meaning that sequence p is the result of appending sequence t to the end of sequence s. The coding is so that if $s = (x)$, the sequence with x as its only element, then $I\Delta_0$ proves $s \leq 9 \cdot x^2$. Morevoer, $I\Delta_0$ proves that for every sequence s, the bound $s \frown (x) \leq 9 \cdot x^2 \cdot s$ holds (see Lemma 3.7 in page 297 of [HP93]). It follows that $I\Delta_0$ proves that if l^r exists below n and s is a sequence of length r all of whose elements are smaller than l, then $s \leq 9^r \cdot l^{2r}$ (by Δ_0-induction on r). Therefore, the coding is fairly close to its information theoretic bound. Of course, $I\Delta_0$ can prove several other obvious facts about $Seq(s)$, $(s)_i$ and $s \frown t$ (see [HP93] for details).

Let $\theta(x, y)$ be a Δ_0-formula in a language L extending the language of arithmetic $\{+, \times, \leq\}$. Obviously, $\theta(x, y)$ defines a binary relation on any infinite model for the language L that may be interpreted as an infinite directed graph. We define Δ_0-formulas $\Theta_i(x, y)$, with certain parameters, meaning that y is reachable from x under certain conditions that depend on the parameters. More precisely, let $\Theta_0(x, y, t, r, l, n) = \theta(x, y)$ (note that the parameters t, r and l are not used for the moment). Inductively, we define $\Theta_{k+1}(x,y,t,r,l,n)$ as follows:

$$(\exists z \le n)(Seq(z) \wedge (z)_0 = x \wedge (z)_r = y \wedge (\forall i < r+1)((z)_i \le l) \wedge \\ \wedge(\forall i < r)(\exists c, c' \le z)((z)_i = c \wedge (z)_{i+1} = c' \wedge \Theta_k(c,c',t,t,l,n))).$$

Informally, the formula $\Theta_k(x,y,t,t,l,n)$ says that y is reachable from x in t^k steps according to the directed graph defined by $\theta(x,y)$ as long as each number in the path is bounded by l. The following theorem states this in the form of recursive equations:

Theorem 22 *Let $M \models I\Delta_0(L)$, and let $x,y,t,r,l,n \in M$ be such that $x,y \le l$, l^t exists in M, $r < t$, and $9^{t+1} \cdot l^{2(t+1)} \le n$. Then,*

(i) $\Theta_{k+1}(x,y,t,0,l,n) \leftrightarrow x = y$,

(ii) $\Theta_{k+1}(x,y,t,r+1,l,n) \leftrightarrow (\exists c \le l)(\Theta_k(x,c,t,t,l,n) \wedge \Theta_{k+1}(c,y,t,r,l,n))$

hold in M.

Proof: (i) Assume $\Theta_{k+1}(x,y,t,0,l,n)$ holds. Then, for some $z \le n$, we have $(z)_0 = x$ and $(z)_0 = y$. Hence, $x = y$. Conversely, if $x = y$, then $z = (x) \le 9 \cdot x^2 \le 9 \cdot l^2 \le n$ is a witness for the existential quantifier in $\Theta_{k+1}(x,y,t,0,l,n)$. (ii) Assume $\Theta_{k+1}(x,y,t,r+1,l,n)$ holds. Let $z \le n$ be the witness for its existential quantifier. Let $c = (z)_1$ and note that $(z)_0 = x$. Obviously, $c \le l$ and $\Theta_k(x,c,t,t,l,n)$ holds. Now let z' be the sequence that results from z when $(z)_0$ is dropped (definable in $I\Delta_0$). It is easily seen that $z' \le z \le n$, and $\Theta_{k+1}(c,y,t,r,l,n)$ holds with z' witnessing its existential quantifier. Conversely, if $\Theta_k(x,c,t,t,l,n)$ and $\Theta_{k+1}(c,y,t,r,l,n)$ hold, let $z' \le n$ be a witness for the existential quantifier in the latter. We can assume that z' codes a sequence of $r+1$ numbers bounded by l each; if not, just trim z' to the first $r+1$ numbers (definable in $I\Delta_0$) and the result is still a witness of the existential quantifier. Moreover, $z' \le 9^{r+1} \cdot l^{2(r+1)}$. Since $x \le l$ and $r < t$, the sequence $(x) \frown z'$ is coded by some $z \le 9^{r+2} \cdot l^{2(r+2)} \le n$. such a z is a witness for the existential quantifier in $\Theta_{k+1}(x,y,t,r+1,l,n)$ and we are done. □

The reader will notice that the recursive equations in Theorem 22 correspond to the inductive definition of the transitive closure of the graph defined by Θ_k. The innermost level of stratification, namely $k = 0$, is the inductive definition of the transitive closure of the graph defined by θ. It is in this sense that we interpret Theorem 22 as a formalization of Nepomnjaščij's Theorem.

5.3 The Proof in Bounded Arithmetic

The graph of the exponentiation function $x = y^z$ is definable by a Δ_0-formula on the natural numbers. Moreover, Pudlák gave a definition with the basic properties being provable in $I\Delta_0$.

Similarly, one can define $y = \lceil \log(x) \rceil$, and $y = \lceil \log^{(k)}(x) \rceil$ in $I\Delta_0$. We make the convention that when expressions such as $\log(a)$ or $(\log(a))^\epsilon$ do not come up integer numbers, the nearest larger integer is assumed unless specified otherwise. Thus, $(\log(a))^\epsilon$ really stands for $\lceil \log(a)^\epsilon \rceil$.

We let L be the usual language of arithmetic $\{+, \times, \leq\}$ extended by a unary function symbol α. We denote $I\Delta_0(L)$ (see the previous section) by $I\Delta_0(\alpha)$. The Weak Pigeonhole Principle PHP_n^m is formalized by the following statement:

$$(\forall x < m)(\alpha(x) < n)) \to (\exists x, y < m)(x \neq y \wedge \alpha(x) = \alpha(y)).$$

We will abbreviate this statement by $\neg\alpha : m \xrightarrow{1-1} n$. We will make use of the following result:

Theorem 23 [PWW88] *For every* $K > 0$,

$$I\Delta_0(\alpha) \vdash (\exists y)(y = x^{\log^{(K)}(x)}) \to \neg\alpha : x^2 \xrightarrow{1-1} x.$$

Here, $\log^{(0)}(x) = x$ *and* $\log^{(k+1)}(x) = \log(\log^{(k)}(x))$.

Next, we formalize the reduction of PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$ using the result of Section 5.2.

Theorem 24 *For every* $K > 0$,

$$I\Delta_0(\alpha) \vdash (\exists y)(y = x^{(\log(x))^{1/K}}) \to \neg\alpha : 2x \xrightarrow{1-1} x.$$

Proof: We will prove that every model of $I\Delta_0(\alpha)$ satisfies the (universal closure of the) formula. The Completeness Theorem for first-order logic will do the rest. Let $\epsilon = 1/K$. Let $\mathbf{M} = (M, F)$ be a model of $I\Delta_0(\alpha)$, let $a \in M$ be such that $a^{(\log(a))^\epsilon}$ exists in M, and assume for contradiction that $F : 2a \xrightarrow{1-1} a$. Let $T = (\log(a))^\epsilon$, $L = a^2$, and $N = 9^{T+1} \cdot L^{2(T+1)}$. Observe that N exists in M since $a^{(\log(a))^\epsilon}$ exists, and M is closed under multiplication. Define $\theta(x, y)$ as follows (see the text that follows the formula for the intuition):

$$(\exists r < 2a)(\exists r' < a)(\exists q, q' < a)(x = 2aq + r \wedge y = aq' + r' \wedge \alpha(r) = r' \wedge q = q').$$

Note that this Δ_0-formula could be informally abbreviated by

$$y \bmod a = \alpha(x \bmod 2a) \wedge \lfloor y/a \rfloor = \lfloor x/2a \rfloor$$

when $x, y \in \{0, \ldots, a^2 - 1\}$.

Lemma 37 $\theta(x, y)^{\mathbf{M}} : 2^{i+1}a \xrightarrow{1-1} 2^i a$ *for every* $i < \log(a)$.

Proof: Given $u \in 2^{i+1}a$, let $v = \lfloor u/2a \rfloor \cdot a + \alpha(u \bmod 2a)$. It is not hard to see that $v \in 2^i a$ and $\theta(u,v)$ holds. Moreover, if $w \in 2^i a$ is such that $\theta(u,w)$ holds, then $\lfloor v/a \rfloor = \lfloor u/2a \rfloor = \lfloor w/a \rfloor$ and $v \bmod a = \alpha(u \bmod 2a) = w \bmod a$. Hence, $v = w$. This shows that $\theta(x,y)^{\mathbf{M}}$ is the graph of a function from $2^{i+1}a$ to $2^i a$.

We show next that the function is one-to-one. Let $u, v \in 2^{i+1}a$ and $w \in 2^i a$ be such that $\theta(u,w)$ and $\theta(v,w)$. Then, $\lfloor u/2a \rfloor = \lfloor v/2a \rfloor = \lfloor w/a \rfloor$ and $F(u \bmod 2a) = F(v \bmod 2a) = w \bmod a$. Since F is one-to-one, it must be then that $u \bmod 2a = v \bmod 2a$. Hence, $u = 2a \cdot \lfloor u/2a \rfloor + (u \bmod 2a) = 2a \cdot \lfloor v/2a \rfloor + (v \bmod 2a) = v$. $\Box$

Lemma 38 $\Theta_K(x,y,T,T,L,N)^{\mathbf{M}} : a^2 \xrightarrow{1-1} a$.

Proof: We prove that for every $k \leq K$, the formula $\Theta_k(x,y,T,T,L,N)$ defines a one-to-one mapping $\Theta_k : 2^{(i+1)T^k}a \to 2^{iT^k}a$ for every $i < \log(a)/T^k$. The lemma will be proved since $T^K = ((\log(a))^{1/K})^K = \log(a)$ (in fact, $T^K \geq \log(a)$ by our convention on rounding). The proof is by induction on k (this induction is outside M).

Lemma 37 takes care of the base case $k = 0$. We turn to the inductive case $0 < k \leq K$. Fix $i < \log(a)/T^k$. We prove that for every $r \leq T$, the formula $\Theta_k(x,y,T,r,L,N)$ defines a one-to-one mapping $\Theta_k^r : 2^{(iT+r)T^{k-1}}a \to 2^{iT^k}a$. That is, we prove that for every $r \leq T$, $x < 2^{(iT+r)T^{k-1}}a$ and $y < 2^{iT^k}a$,

$$\Theta_k(x,z,T,r,L,N) \wedge \Theta_k(y,z,T,r,L,N) \to x = y$$

holds in $\mathbf{M}$. We use the schema of Δ_0-induction on r in the Δ_0-formula above. The base case $r = 0$ is immediate since $\Theta_k(x,y,T,0,L,N)$ defines the identity by Theorem 22. Suppose that $0 < r \leq T$, and that $\Theta_k(x,y,T,r-1,L,N)$ defines a one-to-one mapping $\Theta_k^{r-1} : 2^{(iT+r-1)T^{k-1}}a \to 2^{iT^k}a$. Since $\Theta_{k-1}(x,y,T,T,L,N)$ defines a one-to-one mapping $\Theta_{k-1} : 2^{(iT+r)T^{k-1}}a \to 2^{(iT+r-1)T^{k-1}}a$ by induction hypothesis on k, and since $\Theta_k(x,y,T,r,L,N)$ defines the composition of Θ_{k-1} and Θ_k^{r-1} by Theorem 22 (observe that $x, y \leq L$, $r-1 < T$ and $9^{T+1} \cdot L^{2(T+1)} \leq N$), it follows that $\Theta_k(x,y,T,r,L,N)$ defines a one-to-one mapping $\Theta_k^r : 2^{(iT+r)T^{k-1}}a \to 2^{iT^k}a$ as required. $\Box$

Since Θ_K is a $\Delta_0(\alpha)$ formula, we have that $(M, \Theta_K(x,y,T,T,L,N)^{\mathbf{M}}) \models I\Delta_0(\alpha)$. Moreover, $a^{\log^{(2)}(a)} < a^{(\log(a))^\epsilon}$ exists in M. It follows from Theorem 23 that $\Theta_K(x,y,T,T,L,N)^{\mathbf{M}}$ is not a one-to-one mapping from $a^2 \to a$; a contradiction to Lemma 38. $\Box$

It is well-known that proofs in $I\Delta_0(\alpha)$ translate into bounded-depth LK proofs of polynomial-size. When statements of the form "$f(x)$ exists" are required as in Theorem 22, the translations come up of size $f(n)^{O(1)}$ (see [Kra95], for example). This gives us Theorem 20 as a corollary.

5.4 Infinitely Many Primes

The existence of infinitely many primes is not guaranteed in weak fragments of arithmetic. For example, it is known that I_{open}, Peano Arithmetic with induction restricted to open formulas, has models with a largest prime [MM89]. It is an open problem whether $I\Delta_0$ proves the infinitude of primes. It is known, however, that $I\Delta_0$ augmented with the axiom $(\forall x)(\exists y)(y = x^{\log(x)})$ proves it. In addition, a *single* application of this axiom suffices[1]. More precisely:

Theorem 25 [PWW88] $I\Delta_0 \vdash (\exists y)(y = x^{\log(x)}) \to (\exists y)(y > x \wedge \mathit{prime}(y))$.

The proof of Theorem 25 is quite involved due to the technical difficulties of working within the theory $I\Delta_0$. However, the main idea is based on a fairly standard argument. Before we go on, we proceed to sketch this idea leaving out some of the tedious details of working in $I\Delta_0$. We will try to keep some similarity with the actual proof and point out the crucial points where a delicate argument is required.

Let us assume that there is a maximal prime x and let $y = x^{\log(x)}$. We will derive a contradiction. We start with the observation that for every i such that $1 \leq i \leq y$, the following equality holds

$$\log(i) = \sum_{p \leq x} \{i\}_p \log(p)$$

where the sum ranges over all primes and $\{i\}_p$ is the maximum k such that p^k divides i. On the other hand,

$$\sum_{i=1}^{y} \{i\}_p = \sum_{k \geq 1} \left\lfloor \frac{y}{p^k} \right\rfloor \leq y \sum_{k \geq 1} \frac{1}{p^k} = \frac{y}{p-1}$$

since there are $\lfloor \frac{y}{p^k} \rfloor$ numbers in $\{1, \ldots, y\}$ that are divisible by p^k. Hence,

$$\frac{1}{y} \sum_{i=1}^{y} \log(i) \quad = \quad \frac{1}{y} \sum_{i=1}^{y} \sum_{p \leq x} \{i\}_p \log(p) = \frac{1}{y} \sum_{p \leq x} \log(p) \sum_{i=1}^{y} \{i\}_p \leq \sum_{p \leq x} \frac{\log(p)}{p-1}. \tag{5.10}$$

At this point we would like to derive a large lower bound on the left-hand side of (5.10) and a small upper bound on the right-hand side of (5.10) so that a contradiction is apparent. Carrying out summations of those many terms is not possible, in general, in $I\Delta_0$. However, one could try using the Weak Pigeonhole Principle to prove that there are more numbers that are smaller than the left-hand side than numbers that are bigger than twice the right-hand side. This contradiction would finish the proof since the weak pigeonhole principle is provable. This is the way in which the proof of Paris, Wilkie and Woods proceeds. We will not review it here. However, the need for the large number $x^{\log(x)}$ to exist is made explicit in the following calculations.

[1]This notion of limited use of an axiom also appears in Chapter V, Section 5, Subsection (g) of [HP93].

For a lower bound on the left-hand side of (5.10) we have:

$$\frac{1}{y}\sum_{i=1}^{y}\log(i) \;\geq\; \frac{1}{y}\sum_{i=y/4}^{y}\log(i) \geq \frac{1}{y}\frac{3}{4}y\log(y/4) \geq \frac{3}{4}(\log(x)^2-2),$$

since $y = x^{\log(x)}$. For an upper bound on the right-hand side of (5.10), we split the sum in two:

$$\sum_{p\leq x}\frac{\log(p)}{p-1} \;=\; \sum_{p\leq x^{1-\epsilon}}\frac{\log(p)}{p-1} + \sum_{x^{1-\epsilon}<p\leq x}\frac{\log(p)}{p-1}.$$

Since $x^{1-\epsilon}+1 \leq p \leq x$ in the second summand and there are at most x terms, we can bound this quantity by

$$\sum_{p\leq x^{1-\epsilon}}\frac{\log(p)}{p-1} + x\frac{\log(x)}{x^{1-\epsilon}} \;=\; \sum_{p\leq x^{1-\epsilon}}\frac{\log(p)}{p-1} + x^{\epsilon}\log(x).$$

Repeating, we obtain the upper bound

$$\sum_{p\leq x^{1-2\epsilon}}\frac{\log(p)}{p-1} + x^{\epsilon}\log(x^{1-\epsilon}) + x^{\epsilon}\log(x) \leq \cdots \leq \frac{1}{2}\left(1+\frac{1}{\epsilon}\right)x^{\epsilon}\log(x).$$

Finally setting

$$\epsilon = \frac{1}{\log(x)}$$

we obtain the desired contradiction when x is large since the left-hand side of (5.10) is actually bigger than the right-hand side.

The aim of this section is to show that a weaker axiom suffices, and so the existence of $x^{\log(x)}$ is not the optimal large number assumption. In order to prove that, we will have to resort to the technique of the Weak Pigeonhole Principle mentioned above. However, we are able to reuse it without change from [PWW88], and we obtain our result as a corollary to the results of the previous section.

Theorem 26 $I\Delta_0 \vdash (\exists y)(y = x^{(\log(x))^{1/K}}) \rightarrow (\exists y)(y > x \wedge \mathit{prime}(y))$ *for every* $K > 0$. *Moreover, there exists a model* $M \models I\Delta_0$ *with a non-standard element* $a \in M$ *such that* $a^{(\log(a))^{1/K}}$ *exists in* M *but* $a^{\log(a)}$ *does not.*

Proof: For the second part, let M be a non-standard model of the theory of true arithmetic $\mathrm{Th}(\mathbb{N})$, and let $a \in M$ be a non-standard element. Obviously, $a^{(\log(a))^{1/K}}$ exists in M since the function is total in true arithmetic. Let $N = \{n \in M : (\exists i \in \omega)(M \models n < a^{i(\log(a))^{1/K}})\}$. It is not hard to see that N is a cut of M that is closed under addition and multiplication. It follows that $N \models I\Delta_0$ (see Lemma 5.1.3 in page 64 of [Kra95]). Finally, $a^{(\log(a))^{1/K}}$ still exists in N by absoluteness of the Δ_0-formula expressing the graph of exponentiation. However, $a^{\log(a)}$ does not exist in N because $a^{\log(a)} > a^{i(\log(a))^{1/K}}$ in M for every standard $i \in \omega$.

For the first part, suppose that $a^{(\log(a))^{1/K}}$ exists in $M \models I\Delta_0$. Our goal is to show that no Δ_0-definable function $F : M \to M$ maps $9a\log(a)$ injectively into $8a\log(a)$. The result would follow from Theorem 11 of [PWW88] since then a prime exists in M between a and a^{11}. Let $b = a^{2^K}$ and observe that $b^{(\log(b))^{1/K}} = a^{2^{K+1}(\log(a))^{1/K}}$ exists in M since it is closed under multiplication. By Theorem 24, no Δ_0-definable function maps $2b$ injectively into b. It follows that no Δ_0-definable function maps $\frac{9}{8}b$ injectively into b; otherwise we could compose that function with itself a constant number of times to map $2b$ injectively into b. We conclude that no Δ_0-definable function maps $9a\log(a)$ injectively into $8a\log(a)$; otherwise, we could juxtapose that function with itself to obtain a Δ_0-definable function mapping $\frac{9}{8}b$ injectively into b (break b and $\frac{9}{8}b$ into $a^{2^K-1}/8\log(a)$ blocks of size $8a\log(a)$ and $9a\log(a)$ respectively). □

We note that $I\Delta_0$ proves $(\forall x)(\exists y)(y = x^{\log(x)^\epsilon}) \to (\forall x)(\exists y)(y = x^{\log(x)})$. However, the second part of Theorem 26 implies that $I\Delta_0$ does not prove $(\exists y)(y = x^{\log(x)^\epsilon}) \to (\exists y)(y = x^{\log(x)})$.

Another major open problem in Feasible Number Theory is whether Fermat's Little Theorem is provable in $I\Delta_0$. Berarducci and Intrigila [BI91] point out that one important difficulty is that the modular exponentiation relation $x^y \equiv z \pmod{n}$ is not known to be Δ_0-definable. The situation has changed, however. Very recently, Hesse [Hes01] proved that the modular exponentiation relation on numbers of $O(\log(n))$ bits is first-order definable. A well-known translational argument shows then that $x^y \equiv z \pmod{n}$ is Δ_0-definable. The proof of this result, however, seems to rely on Fermat's Little Theorem, and therefore it is not clear whether the basic properties of modular exponentiation are provable in $I\Delta_0$.

Open Problem 1 *Find a Δ_0 definition of the modular exponentiation relation whose basic properties are provable in $I\Delta_0$; namely, $x^y x^z \equiv x^{y+z} \pmod{n}$ and $(x^y)^z \equiv x^{yz} \pmod{n}$.*

We believe that a positive solution to this open problem would help developping the number theory of $I\Delta_0$ in the same way that the Δ_0-definition of the (non-modular) exponentiation relation helped developping the metamathematics of $I\Delta_0$ [GD82, HP93].

Chapter 6

Conclusions

We view the results of this thesis as contributions to the field of Propositional Proof Complexity in two ways: (1) the proofs of each of our results required the introduction of some new technique to the field that may show useful for later use, and (2) the results resolve important problems that were left open by other researchers of the field and suggest some new conjectures that will help keep the field active. All our results have been published in conferences and appeared in final form in journals, or are submitted for journal publication at the time of writing [AGG01, AGP01, ABE01, Ats01]. The goal of this section is to review and relate all the results of the thesis, and list some of the open problems they suggest.

6.1 The Think Positively Conjecture

We proved that the monotone sequent calculus MLK quasipolynomially simulates LK and every Frege system. This means that although an arbitrary LK proof may use concepts that are not definable by small monotone formulas, one is able to avoid that by a somewhat tricky simulation that distinguishes $n + 1$ cases according to the number of ones in the underlying variables. Recall that the proof is inspired by the fact that, on the class of slice functions, monotone formulas quasipolynomially simulate arbitrary formulas provably with small MLK proofs. This suggest using a similar approach towards collapsing other seemingly different proof systems.

The main open problem in the context of the monotone calculus MLK is whether MLK polynomially simulates LK on monotone sequents, or on contradictory sets of clauses. As we noted in Chapter 3, a negative answer would imply a major breakthrough since then no Frege systems is polynomially bounded. Actually, even more is true. According to Lemma 18, if MLK does not simulate LK polynomially, then the formulas expressing the basic properties of the polynomial-size monotone threshold formulas that arise from Valiant's probabilistic construction, or AKS's explicit construction, would require super-polynomial size in LK.

However, we tend to believe that the actual answer to the problem is positive.

Conjecture 1 *MLK polynomially simulates LK on monotone sequents, even as refutation systems.*

We call this the *Think Positively Conjecture*. The reason to believe in this conjecture is that there ought to be polynomial-size monotone Boolean formulas for the majority function with their basic properties easily provable. That this would suffice is formalized in Lemma 18. A possible approach in trying to prove this would be to consider Valiant's probabilistic construction of the monotone formulas for majority and try to build up a proof of the basic properties by the probabilistic method. Actually, it is easy to see that the only property that we really need is the symmetry of Valiant's formulas.

6.2 Open Problems Related to Res(2)

The main results of Chapter 4 proved that the Weak Pigeonhole Principle and Random CNF formulas require exponential-size refutations in Res(2). The goal was to understand the difficulty in dealing with these tautologies when formulas are allowed to alternate between conjunctions and disjunctions, although in a minimalist way. We showed that the combination of old and new techniques yield exponential lower bounds. Moreover, in view of the known Res($\log$) upper bounds, our results are quite tight in the case of the Weak Pigeonhole Principle. The main open problem is whether our techniques can be used to prove lower bounds for Res(k) for larger k. We believe that the techniques should work, at least, for constant k in the case of Random CNF formulas. However, we have not been able to do that.

Open Problem 2 *Extend the lower bounds of Chapter 4 beyond Res(2).*

As we mentioned in the introduction, it is also known that Resolution cannot polynomially simulate Res(2) [ABE01]. Actually, the separation can be made exponential [AB01], and, in fact, the proof shows that Res(2) requires exponential-size monotone interpolants. It is open, however, whether Res(2) has feasible interpolation, and any answer would have important consequences for the automatizability of Resolution [AB01]. We note that in view of the results of the last section of Chapter 4, it is quite intriguing that the Lovász-Schrijver proof system has feasible interpolation [Pud97].

Open Problem 3 *Investigate the proof system Q. Does it have feasible interpolation?*

Recall that Q polynomially simulates Res(2), and therefore, if it has feasible interpolation so does Res(2).

6.3 Open Problems Related to the WPHP and Primes

We now know that the bound $n^{O(\log n)}$ is not optimal for bounded-depth proofs of WPHP_n^{2n}. The size-depth trade-off proved in Chapter 5 suggests that a general superpolynomial lower bound for all constant depths, if any, should have a quite exotic form.

Conjecture 2 *The Weak Pigeonhole Principle* $WPHP_n^{2n}$ *has polynomial-size proofs in bounded-depth LK.*

One reason to believe in this conjecture is that there ought to be polynomial-size bounded-depth formulas to approximately count with easily derivable properties. The sufficiency of this is not completely obvious, but is somehow well-known. For an explicit statement of this claim, see Pitassi's Thesis [Pit92]. In fact, the computational problem of approximate counting is indeed solvable by uniform bounded-depth families of formulas [Ajt93]. So, the reason to believe in this conjecture is the same as the one we have to believe in the Think Positively Conjecture: the correctness of some algorithm solving a *natural* computational problem should always be provable within the complexity of the problem itself. Obviously, if the solution to Conjecture 2 is positive with a sufficienly uniform proof, we would get a positive solution to Wilkie's problem on whether $I\Delta_0$ proves the infinitude of primes.

Bibliography

[AB87] N. Alon and R. B. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7:1–22, 1987.

[AB01] A. Atserias and M. L. Bonet. On the automatizability of Resolution and related propositional proof systems. Submitted, 2001.

[ABE01] A. Atserias, M. L. Bonet, and J. L. Esteban. Lower bounds for the weak pigeonhole principle and random formulas beyond resolution. Accepted for publication in *Information and Computation*. A preliminary version appeared in ICALP'01, *Lecture Notes in Computer Science* 2076, Springer, pages 1005–1016., 2001.

[ABL01] A. Atserias, M. L. Bonet, and J. Levy. Some observations on Lovász-Schrijver and cutting planes. Manuscript, 2001.

[AGG01] A. Atserias, N. Galesi, and R. Gavaldà. Monotone proofs of the pigeon-hole principle. *Mathematical Logic Quarterly*, 47:461–474, 2001.

[AGP01] A. Atserias, N. Galesi, and P. Pudlák. Monotone simulations of nonmonotone proofs. In *16th IEEE Conference in Computational Complexity*, pages 36–41, 2001. Submitted to *Journal of Computer and System Sciences*.

[Ajt88] M. Ajtai. The complexity of the pigeonhole principle. In *29th Annual IEEE Symposium on Foundations of Computer Science*, pages 346–355, 1988.

[Ajt93] M. Ajtai. Approximate counting with uniform constant depth circuits. In J.-Y Cai, editor, *Advances in Computational Complexity Theory*, DIMACS series in Discrete Mathematics and Theoretical Computer Science, pages 1–20. American Mathematical Society, 1993.

[Ale00] M. Alekhnovich. Mutilated chessboard problem is exponentially hard for resolution. Manuscript, 2000.

[Ats01] A. Atserias. Improved bounds on the weak pigeonhole principle and infinitely many primes from weaker axioms. Accepted for publication in *Theoretical Computer Science*. A prelimiary version appeared in MFCS'01, *Lecture Notes in Computer Science* 2136, Springer-Verlag, pages 148–158., 2001.

[B01́] M. Bílková. Monotone sequent calculus and resolution. *Commentationes Mathematicae Universitatis Carolinae*, 42(3):575–582, 2001.

[BDG+99] M. L. Bonet, C. Domingo, R. Gavaldà, A. Maciel, and T. Pitassi. Non-automatizability of bounded-depth Frege proofs. In *14th IEEE Conference in Computational Complexity*, pages 15–23, 1999.

[BG99] M. L. Bonet and N. Galesi. A study of proof search algorithms for resolution and polynomial calculus. In *40th Annual IEEE Symposium on Foundations of Computer Science*, 1999.

[BI91] A. Berarducci and B. Intrigila. Combinatorial principles in elementary number theory. *Annals of Pure and Applied Logic*, 55:35–50, 1991.

[BKPS98] P. Beame, R. Karp, T. Pitassi, and M. Saks. On the complexity of unsatisfiability proofs for random k-CNF formulas. In *30th Annual ACM Symposium on the Theory of Computing*, pages 561–571, 1998.

[BP96] P. Beame and T. Pitassi. Simplified and improved resolution lower bounds. In *37th Annual IEEE Symposium on Foundations of Computer Science*, pages 274–282, 1996.

[BP97] S. Buss and T. Pitassi. Resolution and the weak pingeonhole principle. In D. Van Dalem and M. Bezem, editors, *Computer Science Logic '96, 10th Annual Conference of the EACSL*, volume 1258 of *Lecture Notes in Computer Science*, pages 149–156. Springer-Verlag, 1997.

[BP01] S. R. Buss and P. Pudlák. On the computational content of intuitionistic propositional proofs. *Annals of Pure and Applied Logic*, 109:49–64, 2001.

[BPR97] M. L. Bonet, T. Pitassi, and R. Raz. Lower bounds for cutting planes proofs with small coefficients. *Journal of Symbolic Logic*, 62(3):708–728, 1997.

[BS90] R. Boppana and M. Sipser. The complexity of finite functions. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A. Elsevier, 1990.

[BSW01] E. Ben-Sasson and A. Wigderson. Short proofs are narrow–resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.

[BT88] S. R. Buss and G. Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62:311–317, 1988.

[Bus95] S. Buss. Some remarks on lengths of proofs. *Archive for Mathematical Logic*, 34:377–394, 1995.

[Bus97] S. R. Buss. Polynomial size proofs of the propositional pigeonhole principle. *Journal of Symbolic Logic*, 52(4):916–927, 1997.

[CCT87] W. Cook, C. R. Coullard, and G. Turán. On the complexity of cutting-plane proofs. *Discrete Applied Mathematics*, 18:25–38, 1987.

[CF90] M. T. Chao and J. Franco. Probabilistic analysis of a generalization of the unit-clause literal selection heuristics. *Information Science*, 51:289–314, 1990.

[Coo71] S. Cook. The complexity of theorem proving procedures. In *3rd Annual ACM Symposium on the Theory of Computing*, pages 151–158, 1971.

[Coo00] S. Cook. The P versus NP problem. Problems of the Millennium. Available at `http://www.claymath.org/prizeproblems/index.htm`, 2000.

[CR79] S. Cook and R. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44:36–50, 1979.

[CR92] V. Chvátal and B. Reed. Mick gets some (the odds are on his side). In *33rd Annual IEEE Symposium on Foundations of Computer Science*, pages 620–627, 1992.

[CS88] V. Chvátal and E. Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, 1988.

[DR00] S. Dantchev and S. Riis. Planar tautologies hard for resolution. Manuscript, 2000.

[ER60] P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hugar. Acad. Sci.*, 5:17–61, 1960.

[For97] L. Fortnow. Time-space tradeoffs for satisfiability. In *12th IEEE Conference in Computational Complexity*, pages 52–60, 1997. To appear in *Journal of Computer and System Sciences*.

[FS90] A. Frieze and S. Suen. Analysis of two simple heuristics on a random instance of k-SAT. *Journal of Algorithms*, 20:312–355, 1990.

[FvM00] L. Fortnow and D. van Meldebeek. Time-space tradeoffs for non-deterministic computation. In *15th IEEE Conference in Computational Complexity*, 2000.

[Gal00] N. Galesi. *On the Complexity of Propositional Proof Systems*. PhD thesis, Universitat Politècnica de Catalunya, 2000.

[GD82] H. Gaifman and C. Dimitracopoulos. Fragments of Peano's arithmetic and the MRDP theorem. In *Logic and algorithmic*, number 30 in Monographies de l'Enseignement Mathématique, pages 187–206. Univeristé de Genève, 1982.

[GS82] G. R. Grimmet and D. R. Stirzaker. *Probability and Random Processes*. Oxford Science Publications, 1982.

[Hak85] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.

[Hes01] W. Hesse. Division is in uniform TC^0. In *28th International Colloquium on Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, pages 104–114. Springer-Verlag, 2001.

[HP93] P. Hájek and P. Pudlák. *Metamathematics of First-Order Arithmetic*. Springer, 1993.

[IPU94] R. Impagliazzo, T. Pitassi, and A. Urquhart. Upper and lower bounds for tree-like cutting planes proofs. In *9th IEEE Symposium on Logic in Computer Science*, pages 220–228, 1994.

[KKKS98] L. M. Kirousis, E. Kranakis, D. Krizanc, and Y. C. Stamatiou. Approximating the unsatisfiability threshold of random formulas. *Random Structures and Algorithms*, 12(3):253–269, 1998.

[KP95] J. Krajícek and P. Pudlák. Some consequences of cryptographical conjectures for s_2^1 and ef. In D. Leivant, editor, *Logic and Computational Complexity*, volume 960 of *Lecture Notes in Computer Science*, pages 210–220. Springer-Verlag, 1995.

[KPT91] J. Krajícek, P. Pudlák, and G. Takeuti. Bounded arithmetic and polynomial hierarchy. *Annals of Pure and Applied Logic*, 52:143–154, 1991.

[KPW95] J. Krajícek, P. Pudlák, and A. Woods. Exponential lower bound to the size of bounded depth Frege proofs of the pigeon hole principle. *Random Structures and Algorithms*, 7(1):15–39, 1995.

[Kra95] J. Krajícek. *Bounded arithmetic, propositional logic, and complexity theory*. Cambridge University Press, 1995.

[Kra00] J. Krajícek. On the weak pigeonhole principle. To appear in *Fundamenta Mathematicæ*, 2000.

[KW90] M. Karchmer and A. Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM Journal on Discrete Mathematics*, 3(2):255–265, 1990.

[LS91] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.

[LV99] R. J. Lipton and A. Viglas. On the complexity of SAT. In *40th Annual IEEE Symposium on Foundations of Computer Science*, pages 459–464, 1999.

[MM89] A. J. Macintyre and D. Marker. Primes and their residue rings in models of open induction. *Annals of Pure and Applied Logic*, 43(1):57–77, 1989.

[MPW00] A. Maciel, T. Pitassi, and A. R. Woods. A new proof of the weak pigeonhole principle. In *32nd Annual ACM Symposium on the Theory of Computing*, 2000.

[Nep70] V. A. Nepomnjaščij. Rudimentary predicates and Turing calculations. *Soviet Math. Dokl.*, 11:1462–1465, 1970.

[PBI93] T. Pitassi, P. Beame, and R. Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3(2):97–140, 1993.

[Pit92] T. Pitassi. *The Power of Weak Formal Systems*. PhD thesis, Department of Computer Science, University of Toronto, 1992.

[Pit97] T. Pitassi. Algebraic propositional proof systems. In N. Immerman and Ph. G. Kolaitis, editors, *Descriptive Complexity and Finite Models*, volume 31 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 68–96. American Mathematical Society, 1997.

[PR01] T. Pitassi and R. Raz. Regular resolution lower bounds for the weak pigeonhole principle. In *33rd Annual ACM Symposium on the Theory of Computing*, pages 347–355, 2001.

[Pud97] P. Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *Journal of Symbolic Logic*, 62(3):981–998, 1997.

[Pud99] P. Pudlák. On the complexity of the propositional calculus. In *Sets and Proofs, Invited Papers from Logic Colloquium '97*, pages 197–218. Cambridge University Press, 1999.

[Pud01] P. Pudlák. On reducibility and symmetry of disjoint NP-pairs. In *26th International Symposium on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, pages 621–632. Springer-Verlag, 2001.

[PWW88] J. B. Paris, A. J. Wilkie, and A. R. Woods. Provability of the pigeonhole principle and the existence of infinitely many primes. *Journal of Symbolic Logic*, 53(4):1235–1244, 1988.

[Raz85a] A. A. Razborov. Lower bounds for the monotone complexity of some boolean functions. *Soviet Math. Doklady*, 31(2):354–357, 1985.

[Raz85b] A. A. Razborov. Lower bounds on monotone complexity of the logical permanent. *Matematicheskie Zametki*, 37(6):887–900, 1985.

[Raz95a] A. A. Razborov. Bounded arithmetic and lower bounds in Boolean complexity. In P. Clote and J. B. Remmel, editors, *Feasible Mathematics II*, volume 13 of *Progress in Computer Science and Applied Logic*, pages 344–386. Birkhäuser, 1995.

[Raz95b] A. A. Razborov. Unprovability of lower bounds on circuit size in certain fragments of bounded arithmetic. Izvestiya of the RAN, 1995.

[Raz96] A. A. Razborov. Lower bounds for propositional proofs and independence results in bounded arithmetic. In *23rd International Colloquium on Automata, Languages and Programming*, volume 1099 of *Lecture Notes in Computer Science*, pages 48–62. Springer-Verlag, 1996.

[Raz98] A. A. Razborov. Lower bounds for the polynomial calculus. *Computational Complexity*, 7:291–324, 1998.

[Raz01a] R. Raz. Resolution lower bounds for the weak pigeonhole principle. Manuscript, 2001.

[Raz01b] A. A. Razborov. Resolution lower bounds for perfect matching principles. Manuscript, 2001.

[RR97] A. A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55:24–35, 1997.

[RW92] R. Raz and A. Wigderson. Monotone circuits for matching require linear depth. *Journal of the ACM*, 39:736–744, 1992.

[RWY97] A. Razborov, A. Wigderson, and A. Yao. Read-once branching programs, rectangular proofs of the pigeonhole principle and the transversal calculus. In *29th Annual ACM Symposium on the Theory of Computing*, pages 739–748, 1997.

[Sha49] C. Shannon. The synthesis of two-terminal switching circuits. *Bell Systems Technical Journal*, 28:59–98, 1949.

[Tak87] G. Takeuti. *Proof Theory*. North-Holland, second edition, 1987.

[Tar87] E. Tardos. The gap between monotone and non-monotone circuit complexity is exponential. *Combinatorica*, 7(4):141–142, 1987.

[Val84] L. G. Valiant. Short monotone formulae for the majority function. *Journal of Algorithms*, 5:363–366, 1984.

[Vol99] H. Vollmer. *Introduction to Circuit Complexity*. Springer-Verlag, 1999.

[Weg87] I. Wegener. *The Complexity of Boolean Functions*. John Wiley & Sons, 1987.

[Woo81] A. R. Woods. *Some problems in logic and number theory, and their connections*. PhD thesis, University of Manchester, Department of Mathematics, 1981.