DD3501 Current Research in Proof Complexity	Dec 01, 2011
Lecture 7	
Lecturer: Jakob Nordström	*Scribe: Mateus de Oliveira Oliveira*

1 Recap of Polynomial Calculus and PCR

In previous lectures, we have seen that in polynomial calculus (PC) a proofs is regarded as a sequence of inference rules that manipulate polynomial equations. In other words, each line in a polynomial calculus derivation is a multivariate polynomial $p \in \mathbb{F}[x, y, z, \ldots]$ for some (fixed) field $\mathbb{F}$ (which we will generally assume to be finite). When translating from CNF formulas to polynomials, a positive literal x is mapped to the expression $\mathrm{tr}(x) = x$ and a negative literal $\overline{x}$ is mapped to the expression $\mathrm{tr}(\overline{x}) = 1 - x$. Furthermore, a clause C is mapped to the expression $\mathrm{tr}(C) = \prod_{z \in C} \mathrm{tr}(z)$ where z runs over the literals of C. Finally, a formula $F = C_1 \wedge C_2 \wedge \cdots \wedge C_n$ is mapped to the system of equations $\mathrm{tr}(C_i) = 0$ for $i = 1, \ldots, n$. In what follows we write simply $\mathrm{tr}(C)$ instead of $\mathrm{tr}(C) = 0$. The rules of inference for the polynomial calculus are as follows:

Boolean axioms $\dfrac{}{x^2 - x}$ for all variables x (forcing 0/1-solutions)

Linear combination $\dfrac{p \qquad q}{\alpha p + \beta q}$

Multiplication $\dfrac{p}{xp}$

where $\alpha, \beta \in \mathbb{F}$, $p, q \in \mathbb{F}[x, y, z, \ldots]$, and x is any variable. As has been noted before, because of the Boolean axioms we can assume without loss of generality that all polynomials appearing in a PC-derivation are multilinear.

In this course we will mostly focus on an extension of polynomial calculus known as polynomial calculus resolution or PCR. Recall that in PCR we introduce new variables $\overline{x}, \overline{y}, \overline{z}, \ldots$, and use the new variable $\overline{x}$ instead of the factor $1 - x$ to encode negation. Note that these new variables are a priori totally independent of everything else. We therefore have to make sure that $\overline{x}$ is the negation of x by adding the following axioms.

Complementarity axioms $\dfrac{}{x + \overline{x} - 1}$ for all variables x.

All other derivation rules and proof complexity measures are as for polynomial calculus.

2 Degree vs Size in Polynomial Calculus

The first goal of this lecture is to prove Theorem 2.1 below, which relates the size of a polynomial calculus proof with the degree of the largest monomial that has to appear in it. This theorem was proven in [IPS99] building on the work [CEI96]. Recall that $S(\pi)$ denotes the size of a derivation (the total number of monomials counted with repetitions) and $Deg(\pi)$ denotes the maximal degree of any monomial appearing in π, and that taking the minimum over all PCR-refutations of an unsatisfiable CNF formula F, $S_{\mathcal{PCR}}(F \vdash \bot)$ and $Deg_{\mathcal{PCR}}(F \vdash \bot)$ denote the minimal size and degree, respectively, of refuting F.

Theorem 2.1 ([IPS99]). *Let F be a unsatisfiable CNF formula of width $W(F)$ over n variables. Then*

$$Deg_{\mathcal{PCR}}(F \vdash \bot) \leq W(F) + \mathrm{O}\big(\sqrt{n \cdot \ln S_{\mathcal{PCR}}(F \vdash \bot)}\big) \ .$$

The bound in Theorem 2.1 is essentially tight, since it was shown in [GL10] that there exist k-CNF formulas F whose refutations require degree at least $Deg_{\mathcal{PCR}}(F \vdash \bot) \approx \sqrt{n}$, but where the refutation size is $S_{\mathcal{PCR}}(F \vdash \bot) = \text{poly}(n)$. We notice that Theorem 2.1 implies that strong lower bounds on the degree of PCR refutations can be translated into strong lower bounds on size as stated in the following corollary.

Corollary 2.2. *Let F be an unsatisfiable k-CNF formula on n variables for $k = O(1)$. Then*

$$S_{\mathcal{PCR}}(F \vdash \bot) = exp\left(\Omega\left(\frac{Deg_{\mathcal{PCR}}(F \vdash \bot)^2}{n}\right)\right) .$$

As a side note, we remark that Theorem 2.1 and Corollary 2.2 were originally stated and proven for PC, but the proofs carry over to PCR essentially verbatim so this is how we will present them here.

Since a clause of width w gets translated to a PCR-monomial of degree w, and in view of the similarity between Theorem 2.1 and the bound

$$L_{\mathcal{R}}(F \vdash \bot) = \exp\left(\Omega\left(\frac{(W_{\mathcal{R}}(F \vdash \bot) - W(F))^2}{n}\right)\right) \tag{2.1}$$

relating length and width in resolution [BW01], it is natural to consider width in resolution as in some sense analogous to degree in PCR. In the same way, since clauses correspond to monomials, clause space in resolution is an analogous measure to monomial space in PCR (as suggested by the notation we have adopted). It is a natural (and interesting) question how far these analogies go. In particular, is there any analogue of [AD08] showing that monomial space is an upper bound on degree? This question is wide open as far as the lecturer is aware.

Open Problem 1. *Is there an upper bound on refutation degree in terms of refutation monomial space in PCR similar to what holds for width with respect to clause space in resolution by [AD08]?*

Before starting to prove Theorem 2.1, we need to take care of a couple of technicalities. First we notice that even though PCR may be more efficient than PC with respect to monomial space, the degree measure is the same for the two systems.

Proposition 2.3. *For any CNF formula F it holds that $Deg_{\mathcal{PC}}(F \vdash \bot) = Deg_{\mathcal{PCR}}(F \vdash \bot)$.*

Proof. Left as an easy exercise (but some care is needed to get the formal details right since the encodings of the formula F in the two proof systems is different). □

What this means is that a strong enough degree lower bound in either PC or PCR is enough to prove asymptotically the same lower bound on proof size in both proof systems at the same time.

We also notice that just as for resolution, restrictions preserve refutations in both PC and PCR. Again we leave the proof to the reader as a straightforward exercise.

Proposition 2.4. *If $\pi : F \vdash \bot$ is a PC- or PCR-refutation, and ρ is a restriction of the variables in F, then $\pi\restriction_\rho$ is a refutation of $F\restriction_\rho$ in at most the same size, length, space and degree as π.*

We now turn to the proof of Theorem 2.1. Given an unsatisfiable CNF formula F of width $k = W(F)$ with n variables, and requiring proofs of size $S = S_{\mathcal{PCR}}(F \vdash \bot)$ to be refuted, we set

$$d = \sqrt{2n \ln S} \tag{2.2}$$

and

$$a = \left(1 - \frac{d}{2n}\right)^{-1} \tag{2.3}$$

(where we note that we can prove that $S < \exp(2n)$ by doing a good enough simulation of resolution in PCR, so $d < 2n$ and hence $a > 1$). Let $\mathit{fat}(\pi)$ denote the number of monomials of degree at least d ("fat monomials") in a proof π of the unsatisfiability of F. We will do the proof of Theorem 2.1 by induction on the magnitude of $\mathit{fat}(\pi)$: If $\mathit{fat}(\pi)$ is zero, then we are done; otherwise we will apply an inductive argument which is implicit in the following lemma.

Lemma 2.5. *Let G be a k-CNF formula over $m \leq n$ vars. Suppose, using the definitions in Equations (2.2) and (2.3), that there exists a PCR-refutation $\pi : G \vdash \perp$ such that $\mathit{fat}(\pi) < a^b$ for some $b \in \mathbb{N}$. Then there exists a PCR-refutation π' of G such that $\mathit{Deg}(\pi') \leq k + d + b - 1$.*

Deferring the proof of Lemma 2.5 for a moment, let us see how Theorem 2.1 immediately follows from this lemma.

Proof of Theorem 2.1. Set $G = F$ and $m = n$ and apply Lemma 2.5. Not all monomials in a refutation π of F can be fat (for instance, the final contradiction 1 is not, and neither are the initial axiom clauses—viewed as monomials—if the formula has constant width), so for any refutation π—including the smallest one there is of size $S = S_{\mathcal{PCR}}(F \vdash \perp)$—we have

$$\mathit{fat}(\pi) < S < a^{\lceil \log_a S \rceil} \leq a^{1 + \lfloor \log_a S \rfloor} \ . \tag{2.4}$$

Let us set

$$b = 1 + \lfloor \log_a S \rfloor \ . \tag{2.5}$$

Then $\mathit{fat}(\pi) < a^b$, so Lemma 2.5 says that there exists a PCR-refutation π' with $\mathit{Deg}(\pi') \leq k + d + b - 1$. The k in this inequality is what we have in Theorem 2.1, and $d = \sqrt{2n \ln S}$ also has the right order of magnitude. It remains to take a closer look at b.

By using the equality $\log_a x = \ln x / \ln a$ we see that

$$1 + \log_a S = 1 + \frac{\ln S}{\ln a} \ . \tag{2.6}$$

Furthermore, by the definition of a we get for the denominator in (2.6) that

$$\ln a = \ln \left(1 - \frac{d}{2n}\right)^{-1} = -\ln \left(1 - \frac{d}{2n}\right) \geq d/2n \ , \tag{2.7}$$

where we used the standard inequality $\ln(1 + x) \leq x$ (valid for all $x > -1$), which implies that $-\ln(1 - x) \geq x$. Combining Equations (2.6) and (2.7) and plugging them into Equation (2.5), we obtain

$$b \leq 1 + \log_a S \leq 1 + \sqrt{2n \ln S} \ . \tag{2.8}$$

Thus, assuming that Lemma 2.5 is true, we have that

$$\mathit{Deg}(\pi') \leq k + d + b - 1 \leq W(F) + 2\sqrt{2n \log S} \tag{2.9}$$

and the theorem follows. □

Proof of Lemma 2.5. The proof is by nested induction on b and on the number of variables $m \leq n$ in G. Let us first consider some base cases:

1. If $b = 0$ we are done, since there are no fat monomials (that is, the PCR refutation already has sufficiently small degree).

2. If $S = 1$ then the formula must contain the empty clause, since there is not enough space to even perform an inference step. Just listing the monomial 1 that is the translation of the empty clause is a refutation in degree 0.

3. $m = 1$ or $k = 1$: In this case G must contain $x \wedge \overline{x}$ for some variable x. If $m = 1$ then G has exactly one variable that must be resolved. If the width of G is $k = 1$, then G also must have complementary unit clauses. The resolution refutation that resolves x and $\overline{x}$ can clearly be simulated in PCR in degree 1. Since we have to perform at least one addition to do a (nontrivial) PCR refutation, we have $S \geq 4$ and $d = \sqrt{2n \ln S} > 1$ and the degree bound in the lemma holds.

4. $m \leq k$: Again we observe that $S \geq 4$, implying that $d = \sqrt{2n \ln S} > 1$. Also, the degree of the refutation cannot be greater than the number of variables. Thus $Deg(\pi') \leq m \leq k \leq k + d + b - 1$, since $d + b - 1 \geq 0$.

Now for the inductive step. Let $\pi : G \vdash \bot$ be a refutation of G and let $fat(\pi) < a^b$. Then there are at most $2m$ distinct literals and at least $d \cdot fat(\pi)$ literals in fat monomials. By counting, there exists a literal (say, x) that appears in at least $d \cdot fat(\pi)/2m \geq d \cdot fat(\pi)/2n$ monomials. Set this literal x to true (which we agreed is 0 in our PC/PCR-universe), making in this way all the monomials containing x vanish. Then $\pi\restriction_x : G\restriction_x \vdash \bot$ has strictly less than $(1 - d/2n) \cdot a^b \leq a^{b-1}$ fat monomials. It follows from the induction hypothesis that

$$Deg_{\mathcal{PCR}}(G\restriction_x \vdash \bot) \leq k + d + b - 2 \ . \tag{2.10}$$

Now consider $\pi\restriction_{\overline{x}} : G\restriction_{\overline{x}} \vdash \bot$. Here we do not have any guarantees that many fat monomials disappear, but since the number of variables in $G\restriction_{\overline{x}}$ is one less we get from the induction hypothesis that

$$Deg_{\mathcal{PCR}}(G\restriction_{\overline{x}} \vdash \bot) \leq k + d + b - 1 \ . \tag{2.11}$$

Now if we take a refutation of $G\restriction_x$ with degree as in (2.10) and multiply by $\overline{x}$ everywhere, we obtain a derivation of $\overline{x}$ from G in degree at most $k + d + b - 1$. Subtracting $x + \overline{x} - 1$ from $\overline{x}$ we get $1 - x$, which we can then use to derive $G\restriction_{\overline{x}}$ from G without increasing the degree. (For any clause $C \vee x$ in G translated to the monomial $m(C)x$, multiply $1 - x$ by $m(C)$ and then subtract to obtain $m(C)$ encoding the restricted clause C.) Finally, append a refutation of $G\restriction_{\overline{x}}$ of degree as in (2.11). This proves the inductive step, and the lemma follows. □

3 (Towards) Exponential Lower Bounds on Proof Size in PCR

Now we want to use Theorem 2.1 to prove strong lower bounds on PCR (and PC) size. We will also use the opportunity to get acquainted with a new formula family.

3.1 Random k-CNF Formulas

Definition 3.1. F is a *random k-CNF formula* with Δn clauses over n variables, denoted $F \sim \mathcal{F}_k^{n,\Delta}$, if F is constructed by picking Δn clauses independently and uniformly at random with replacements from the set of all $\binom{n}{k} \cdot 2^k$ k-clauses over n variables.

We say that Δ is the clause density of F. In these notes, we will focus on constant clause density, although in general Δ can be allowed to grow as a function of n.

Intuitively, if Δ is small, then there are not many constraints to be satisfied and in this way, the formula is likely to be satisfiable. As Δ increases the formulas should become harder and harder to satisfy, and at some point they should be very likely to be unsatisfiable. This intuition is correct. In what follows, we will say that something holds *almost surely* if the probability of the event in question is $1 - o(1)$, i.e., if it approaches 1 as n goes to infinity.

Fact 3.2. *For any k, if we pick $\Delta = \mathrm{O}(1)$ large enough (depending on k) it holds for $F \sim \mathcal{F}_k^{n,\Delta}$ that F is almost surely unsatisfiable.*

It is believed that for any k there is a sharp threshold Δ_k such that for any $\epsilon > 0$

- $F \sim \mathcal{F}_k^{n,\Delta_k-\epsilon}$ is almost surely satisfiable,
- $F \sim \mathcal{F}_k^{n,\Delta_k+\epsilon}$ is almost surely unsatisfiable.

For $k = 2$, such a threshold is known to exist and is $\Delta_2 = 1$ (two references for this result are [CR92, Goe96]). For $k = 3$, Δ_3 is believed to be roughly 4.2, but there is no formal proof of this fact or even of that such a sharp threshold does exist. We do not quite have time to go into details about this, but what can briefly be said is that there is some empirical as well as theoretical evidence for a sharp threshold, and also some non-rigorous methods in statistical physics that claim to prove this threshold (but these proofs are not what we usually mean by proofs in the mathematical sense). Regardless of whether such a threshold exists, one can determine a lower bound on the density for which random 3-CNF formulas are almost surely satisfiable and an upper bound above which formulas are almost surely unsatsfiable:

- $F \sim \mathcal{F}_3^{n,\Delta}$ is almost surely satisfiable if $\Delta \leq 3.5$ (see [KKL03]).
- $F \sim \mathcal{F}_3^{n,\Delta}$ is almost surely unsatisfiable if $\Delta \geq 4.6$ (see, for instance, [KKKS98, JSV00, KKS$^+$07]).

There is a whole industry improving these bounds by ingenious methods, not seldom in the third decimal place or so. Since this is not what we are focusing on right now, the results cited above have simply been rounded (in the right direction) to one decimal place (*so note, in particular, that the results as stated above are not the optimal bounds reported in those papers*).

Random k-CNF formulas of small but sufficiently large constant density are plausible candidates to be hard for pretty much any proof system. The intuition for this is that they lack any apparent structure or "reason" for being unsatisfiable (other than the density being large). Unfortunately, this lack of structure also makes it difficult to prove lower bounds for such CNF formulas. They are known to be hard for resolution with respect to both length [CS88] and clause space [BG03]. They are also hard for polynomial calculus (both PC and PCR) with respect to size. This lower bound was first proven by Ben-Sasson and Impagliazzo [BI99] for fields of characteristics distinct from 2, and subsequently improved to fields of any characteristics including 2 (which is the most important case for us) by Alekhnovich and Razborov [AR01], the journal version of which is [AR03]. This family of formulas is not known be hard for cutting planes however, and this is arguably a major open problem in the field.

Open Problem 2. *Prove that random k-CNF formulas are hard for Cutting Planes with respect to proof size.*

A less major, but still interesting (and annoying) problem is to show that random formulas have large space complexity for PCR

Open Problem 3. *Prove that random k-CNF formulas are hard for PCR (or even for PC) with respect to monomial space.*

As things stand now, even proving a lower bound for PC would be great, since we cannot rule out the possibility that random k-CNF formulas can be refuted (with decent probability) even in constant monomial space and even in PC (although this would seem slightly absurd).

Of course, these are both open problems, so strictly speaking it could be that there are no lower bounds to be proven. This would be very surprising, however, or at least this is what the lecturer thinks.

3.2 Linear Lower Bound on Refutation Degree for Random k-CNF Formulas

Our goal for this and (mainly) the next lecture is to prove a result in [AR03] (one of many in that paper) establishing a linear lower bound on the degree of PC/PCR refutations of random k-CNF formulas. Combining such a lower bound with Corollary 2.2, we can obtain an exponential lower bound on the size of PC/PCR refutations. To show the degree lower bound it is sufficient for us to focus on PC, since by Proposition 2.3 the same lower bound will then apply to PCR as well.

Recall that a bipartite graph $G = (U \cup V, E)$ is a (d, s, e)-unique-neighbour expander, or boundary expander, if

1. Every vertex in U has at most d incident edges.
2. For all $U' \subseteq U$ s.t. $|U'| \leq s$, it holds that $|\partial U'| \geq e \cdot |U'|$, where the boundary $\partial U'$ is defined as $\partial U' = \{v \in N(u')\,;\, |N(v) \cap U'| = 1\}$.

Remark 3.3. To readers of [AR03], we should perhaps mention that our presentation will be slightly different on a syntactical level. In [AR03], the argument is made in terms of so-called (r', s', c')-*expander matrices*, but these are the same as bipartite boundary expanders if we identify the left-hand side vertices U with the rows and the right-hand side vertices V with the columns of the matrix, and let s' be the degree d, r' the maximum size of the expanding set s, and c' the expansion constant e.

Definition 3.4 (Graph representation of F). For F a CNF formula we define the *graph representation* of F to be the bipartite graph $G(F) = (U \cup V, E)$ where each vertex in U corresponds to a clause of F, each vertex in V corresponds to a variable in $Vars(F)$, and there is an edge $(C, x) \in E$ precisely if the variable x appears (positively or negatively) in the clause C.

Lemma 3.5 ([CS88]). *Let $F \sim \mathcal{F}_k^{n,\Delta}$ for $\Delta = \mathrm{O}(1)$, and let $G(F)$ be the bipartite graph representing the formula F. Then there are constants $\kappa_1, \kappa_2 > 0$ such that $G(F)$ is almost surely a $(k, \kappa_1 \cdot n, \kappa_2)$-boundary expander.*

Recall that we had a very similar claim when we studied the formulas $PHP(G)$ obtained from random restrictions of pigeonhole principle formulas and proved lower bounds for them in lecture 3. Working out the details of this former claim was a problem on the first problem set, and since the proof of Lemma 3.5 is similar in spirit we will just accept it as true. We remark that [AR03] states a more general theorem for non-constant Δ, but we focus on $\Delta = O(1)$. Alekhnovich and Razborov prove that if we take a formula F for which $G(F)$ is a good expander, then this expansion property is sufficient to guarantee a strong lower bound on PCR refutation degree.

Theorem 3.6 ([AR03]). *If $G(F)$ is a (d, s, e)-boundary expander, then $Deg_{\mathcal{PC}}(F \vdash \perp) \geq se/2$ for any field $\mathbb{F}$.*

This is the theorem that we want to prove in the next lecture. Combining this with Corollary 2.2, we have the following result.

Corollary 3.7 ([AR03]). *If $F \sim \mathcal{F}_k^{n,\Delta}$ for $k \geq 3$ and $\Delta = \mathrm{O}(1)$ large enough, then almost surely F is unsatisfiable and $S_{\mathcal{PCR}}(F \vdash \perp) = \exp(\Omega(n))$ over any field $\mathbb{F}$.*

The reason we need to pick Δ large is that we want F to be unsatisfiable.

3.3 Setting Up Some Notation and Terminology

Recall that a CNF Formula $F = C_1 \wedge .... \wedge C_m$ for $x_1, ..., x_n$ is represented in PC by polynomials $f_1, ..., f_m$ in $\mathbb{F}[x_1, ..., x_n]$, one polynomial per clause. We may regard a PC-derivation as computing elements in the ideal in $\mathbb{F}[x_1, ..., x_n]$ generated by $f_1, ..., f_m$, and a refutation can be regarded as the process of showing that 1 belongs to this ideal. Let us now introduce some notation that will be used in the next lecture to prove Theorem 3.6:

- T_n is the set of all multilinear monomials or terms over variables $x_1, \ldots, x_n$.
- $Deg(m)$ is the degree of the monomial/term m, which is equal to the number of variables (since m is multilinear).
- $T_{n,d} = \{t \in T_n \mid Deg(t) \leq d\}$ is the set of all (multilinear) monomials of degree at most d.
- $S_n(\mathbb{F})$ is the set of all (multilinear) polynomials with coefficients in $\mathbb{F}$ (where we recall that the multilinearity restriction is without loss of generality due to the Boolean axioms).
- $S_{n,d}$ is the linear space of all multilinear polynomials of degree at most d.

Again, as noted at the very beginning of this lecture we are implicitly assuming that the Boolean axioms $x^2 - x$ are always applied to get multilinear polynomials. If we wanted to be more formal, we could have said that we are working in the polynomial ring $\mathbb{F}[x_1, ..., x_n]$ factored by the ideal I generated by $\{x_i^2 - x_i \mid i = 1, \ldots, n\}$, which is an $\mathbb{F}$-algebra denoted by $S_n(\mathbb{F})$ in [AR03]. We will use the same notation, but sort of gloss over these algebraic niceties.

References

[AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, May 2008. Preliminary version appeared in *CCC '03*.

[AR01] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS '01)*, pages 190–199, October 2001.

[AR03] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Available at http://people.cs.uchicago.edu/~razborov/files/misha.pdf. Preliminary version appeared in *FOCS '01*.

[BG03] Eli Ben-Sasson and Nicola Galesi. Space complexity of random formulae in resolution. *Random Structures and Algorithms*, 23(1):92–109, August 2003. Preliminary version appeared in *CCC '01*.

[BI99] Eli Ben-Sasson and Russell Impagliazzo. Random CNF's are hard for the polynomial calculus. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS '99)*, pages 415–421, October 1999.

[BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version appeared in *STOC '99*.

[CEI96] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996.

[CR92] Vašek Chvátal and Bruce A. Reed. Mick gets some (the odds are on his side). In *Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science (FOCS '92)*, pages 620–627, October 1992.

[CS88] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988.

[GL10] Nicola Galesi and Massimo Lauria. Optimality of size-degree trade-offs for polynomial calculus. *ACM Transactions on Computational Logic*, 12:4:1–4:22, November 2010.

[Goe96] Andreas Goerdt. A threshold for unsatisfiability. *Journal of Computer and System Sciences*, 53(3):469–486, December 1996.

[IPS99] Russell Impagliazzo, Pavel Pudlák, and Jiri Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999.

[JSV00] Svante Janson, Yannis C. Stamatiou, and Malvina Vamvakari. Bounding the unsatisfiability threshold of random 3-SAT. *Random Structures and Algorithms*, 17(2):103–116, 2000.

[KKKS98] Lefteris M. Kirousis, Evangelos Kranakis, Danny Krizanc, and Yannis C. Stamatiou. Approximating the unsatisfiability threshold of random formulas. *Random Structures and Algorithms*, 12(3):253–269, May 1998.

[KKL03] Alexis C. Kaporis, Lefteris M. Kirousis, and Efthimios G. Lalas. Selecting complementary pairs of literals. *Electronic Notes in Discrete Mathematics*, 16:47–70, October 2003.

[KKS+07] Alexis C. Kaporis, Lefteris M. Kirousis, Yannis C. Stamatiou, Malvina Vamvakari, and Michele Zito. The unsatisfiability threshold revisited. *Discrete Applied Mathematics*, 155:1525–1538, June 2007.