

Locally consistent constraint satisfaction problems

Zdeněk Dvořák Daniel Král' Ondřej Pangrác *

Abstract

An instance of a constraint satisfaction problem is l -consistent if any l constraints of it can be simultaneously satisfied. For a set Π of constraint types, $\rho_l(\Pi)$ denotes the largest ratio of constraints which can be satisfied in any l -consistent instance composed by constraints from the set Π . In the general case of sets Π consisting of finitely many Boolean predicates, we express the limit $\rho_\infty(\Pi) := \lim_{l \rightarrow \infty} \rho_l(\Pi)$ as the minimum of a certain functional on a convex set of polynomials. Our technique yields a robust deterministic algorithm (for a fixed set Π) running in time linear in the size of the input and $1/\varepsilon$ which finds either an inconsistent set of constraints (of size bounded by the function of ε) or a truth assignment which satisfies the fraction of at least $\rho_\infty(\Pi) - \varepsilon$ of the given constraints. In addition, we compute the values of $\rho_l(\{P\})$ for every predicate P which has arity at most three or which is 1-extendable.

1 Introduction

Constraint satisfaction problems form an important abstract computational model for a lot of problems arising in practice. This is witnessed by an enormous recent interest in the computational complexity of various constraint satisfaction problems [4, 5, 6, 18]. Some instances of real problems do not

*Department of Applied Mathematics and Institute for Theoretical Computer Science (ITI), Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic, E-mail: {rakdver,kral,pangrac}@kam.mff.cuni.cz. Institute for Theoretical Computer Science is supported by Ministry of Education of Czech Republic as project LN00A056.

require all the constraints to be satisfied but it is enough to satisfy a large fraction of them. In order to maximize the fraction of satisfied constraints, the input can usually be pruned by removing small sets of contradictory constraints so the input instance is “locally” consistent. Formally, an instance of the constraint satisfaction problem is *l-consistent* if any l constraints of it can be simultaneously satisfied.

In this paper we focus on constraint satisfaction problems whose constraints are Boolean predicates. If Π is a set of Boolean predicates, then $\rho_l(\Pi)$ denotes the fraction of the constraints which can be satisfied in each l -consistent instance of the problem whose constraints are the predicates of Π . If Π consists of only a single predicate P , we simply use $\rho_l(P)$ instead of $\rho_l(\{P\})$. Similarly, $\rho_l^w(\Pi)$ denotes this ratio for the weighted version of the problem (see Section 2 for formal definitions). The limits of $\rho_l(\Pi)$ and $\rho_l^w(\Pi)$ are denoted by $\rho_\infty(\Pi)$ and $\rho_\infty^w(\Pi)$, i.e., $\rho_\infty(\Pi) = \lim_{l \rightarrow \infty} \rho_l(\Pi)$ and $\rho_\infty^w(\Pi) = \lim_{l \rightarrow \infty} \rho_l^w(\Pi)$.

We study both the asymptotic behavior of $\rho_l(\Pi)$ for finite sets of predicates Π and the exact values of $\rho_l(P)$ for some predicates P . In the following subsections, we discuss our results in more detail as well as their relation to previous work. Most of our results hold both for the weighted and unweighted case. Some of the results even extend to the case when the set Π is infinite (see Section 7).

1.1 Asymptotic results

We express $\rho_\infty^w(\Pi)$ for all finite sets of predicates Π and $\rho_\infty(\Pi)$ for all such sets of predicates Π of arities at least two as the minimum of a certain functional Ψ on a convex hull of a finite set $\pi(\Pi)$ of polynomials derived from Π (Corollary 2). The formal definitions of the functional Ψ and the set $\pi(\Pi)$ are postponed to Section 2.

One of our algorithmic results (Theorem 1) is designing, for any fixed set Π of Boolean predicates, a deterministic algorithm which given $\varepsilon > 0$ and a sufficiently locally consistent instance of the weighted constraint satisfaction problem with total weight w_0 finds a truth assignment which satisfies the constraints whose weight is at least $(\rho_\infty^w(\Pi) - \varepsilon)w_0$. The running time of the algorithm is, for a fixed set Π , linear in the number of the input constraints and $1/\varepsilon$. The algorithm is robust in the sense that if it fails to find the desired truth assignment, then it outputs an inconsistent set of constraints contained in the input whose size is bounded by the function of ε . However,

$\sigma(P)$	P	$l = 1$	$l \geq 2$
1	x	1/2	1
1	$x \wedge y$	1/4	1
2	$x \Leftrightarrow y$		1/2
3	$x \vee y$		3/4

Table 1: The values $\rho_l(P)$ for all non-isomorphic essentially unary and binary Boolean predicates.

it might find a good truth assignment even if the input instance is not sufficiently locally consistent (in particular, the algorithm does not determine the local consistency of the input instance). Finally, the presented algorithm is asymptotically optimal in the sense that the ratio of the weights of satisfied constraints can be made arbitrarily close to $\rho_\infty^w(\Pi)$ by choosing the input parameter ε to be sufficiently small.

1.2 Single-element sets Π

We determine the values of $\rho_l(P)$ for every $l \geq 1$ and every Boolean predicate P which has arity at most three (see Tables 1 and 2) or which is 1-extendable. A predicate P is said to be *1-extendable* if it has the following property: If we fix one of its arguments, we can choose the remaining ones in such a way that the predicate is satisfied. In particular, the 0-ary Boolean predicate which is constantly true is 1-extendable. Let us remark that, in this case, all our results hold both for the unweighted and weighted versions of the studied problems, i.e., the instances witnessing the upper bounds contain each constraint at most once and our lower bound proofs translate smoothly for instances with weighted constraints. From the algorithmic point of view, our results can be interpreted in the following way: The simplest probabilistic algorithms (of the kind used in [11, 16, 19]) are approximation algorithms for locally consistent CSPs with optimum worst-case performance.

Let us point out a somewhat exceptional case of the predicate $P(x, y, z) = x \wedge (y \vee z)$ which is not 1-extendable (fix x to be false). Therefore, our general Theorem 3 does not apply. In Section 6, we show that the values $\rho_l(P)$ are closely related to the corresponding values for 2-SATs. In order to analyze this predicate, we develop a formula for the case of locally consistent 2-SATs in Section 5 (Corollary 3).

$\sigma(P)$	P	$l = 1$	$l = 2$	$l = 3$	$l = 4$	$l = 5$	$l \geq 6$	$l \rightarrow \infty$
1	$x \wedge y \wedge z$	1/8	1	1	1	1	1	1
2	$x \Leftrightarrow y \Leftrightarrow z$	1/4	1/4	1/4	1/4	1/4	1/4	1/4
	$x \wedge (y \Leftrightarrow z)$	1/4	8/27	1/2	1/2	1/2	1/2	1/2
3	exactly one	3/8	3/8	3/8	3/8	3/8	3/8	3/8
	$x \wedge (y \vee z)$	3/8	$\frac{2\sqrt{3}}{9}$	1/2	$\frac{\sqrt{5}-1}{2}$	2/3	$\rho_{l-2}(\Pi_2\text{-SAT})$	3/4
	$(x \Leftrightarrow y) \wedge (x \Rightarrow z)$	3/8	3/8	3/8	3/8	3/8	3/8	3/8
4	$x \Rightarrow y \Rightarrow z$	1/2	1/2	1/2	1/2	1/2	1/2	1/2
	$(x \wedge y) \Leftrightarrow z$	1/2	1/2	1/2	1/2	1/2	1/2	1/2
	at most one	1/2	1/2	1/2	1/2	1/2	1/2	1/2
	one or three	1/2	1/2	1/2	1/2	1/2	1/2	1/2
5	\neg exactly one	5/8	5/8	5/8	5/8	5/8	5/8	5/8
	$x \vee (y \wedge z)$	5/8	5/8	5/8	5/8	5/8	5/8	5/8
	$(x \Leftrightarrow y) \vee (x \wedge z)$	5/8	5/8	5/8	5/8	5/8	5/8	5/8
6	$\neg(x \Leftrightarrow y \Leftrightarrow z)$	3/4	3/4	3/4	3/4	3/4	3/4	3/4
	$x \vee (y \Leftrightarrow z)$	3/4	3/4	3/4	3/4	3/4	3/4	3/4
7	$x \vee y \vee z$	7/8	7/8	7/8	7/8	7/8	7/8	7/8

Table 2: The values $\rho_l(P)$ for all non-isomorphic essentially ternary Boolean predicates.

1.3 Previous work

Constraint satisfaction problems whose constraints are Boolean predicates can be traced back to the late 1970's. Schaefer [15] proved that the decision problem whether a given set of predicates (with allowed negations in their arguments) from a set Π is satisfiable is NP-complete unless each predicate of Π can be defined by a CNF formula consisting only of clauses of size at most two or each predicate of Π can be described by a system of linear equations, i.e., the truth assignment which satisfies it form an affine subspace over $\text{GF}(2)$. However, even if the decision problem can be solved in a polynomial time, the problem to maximize the number of satisfied predicates can still be hard, e.g., Håstad [8] showed that there is no $(2 - \varepsilon)$ -approximation algorithm for the case when the set Π contains a single predicate $P(x_1, x_2, x_3) = (x_1 + x_2 + x_3) \bmod 2$ unless $P = NP$. Note that $\rho_l(\Pi) = 1/2$ for every $l \geq 1$ in this case. In particular, $\rho_\infty(\Pi) = 1/2$ and the algorithm from Theorem 1 achieves the best possible ratio.

One of the most studied variant of the constraint satisfaction problem are locally consistent CNF formulas in which the clauses of a formula are viewed as the given constraints. The corresponding set Π_{SAT} of the predicates is just the set of all the disjunctions. Similarly, $\Pi_{2\text{-SAT}}$ denotes the set $\{(x_1), (x_1 \vee x_2)\}$ of the predicates corresponding to clauses of a 2-SAT formula. The interest in this case is witnessed by a separate section (20.6) devoted to this concept in a recent monograph on extremal combinatorics by Jukna [10]. The exact values of $\rho_l^w(\Pi_{\text{SAT}})$ and $\rho_l^w(\Pi_{2\text{-SAT}})$ are known only for small values of l : clearly, $\rho_1^w(\Pi_{\text{SAT}}) = \rho_1^w(\Pi_{2\text{-SAT}}) = 1/2$. Lieberherr and Specker [12] showed that $\rho_2^w(\Pi_{\text{SAT}}) = \rho_2^w(\Pi_{2\text{-SAT}}) = \frac{\sqrt{5}-1}{2} \approx 0.6180$ and subsequently [13] they showed that $\rho_3^w(\Pi_{\text{SAT}}) = \rho_3^w(\Pi_{2\text{-SAT}}) = 2/3$. Later, these proofs have been simplified by Yannakakis [19] using a probabilistic argument. The case of 4-locally consistent CNF formulas somewhat surprisingly differs from the previous ones: First, $\rho_4^w(\Pi_{\text{SAT}}) \approx 0.6992$ but $\rho_4^w(\Pi_{2\text{-SAT}}) > 0.6992$. Second, the values $\rho_l^w(\Pi_{\text{SAT}})$ for $l = 1, 2, 3$ coincide with the corresponding values defined for a ‘‘fractional’’ version of the problem (which are known for all $l \geq 1$ [11] and are equal to so-called Usiskin's numbers [17]) but the value $\rho_4^w(\Pi_{\text{SAT}})$ differs from the value 0.6920 for the fractional version of the problem.

The asymptotic behavior of $\rho_l^w(\Pi_{\text{SAT}})$ was first addressed by Huang and Lieberherr [9] who showed that $\rho_\infty^w(\Pi_{\text{SAT}}) \leq 3/4$. The limit was settled by Trevisan [16] who showed $\rho_\infty^w(\Pi_{\text{SAT}}) = 3/4$. Trevisan's result also yields that

$\rho_\infty^w(\Pi_{\text{SAT}}) = 3/4$. The latter result can be easily derived from our general expression for $\rho_\infty^w(\Pi)$ as demonstrated in Examples 1 and 3.

2 Notation and preliminaries

Throughout this paper, we only deal with constraints which are Boolean predicates and so we prefer to call them *predicates* to emphasize their kind. For a fixed set Π of (types of) Boolean predicates, let Σ be a set of predicates whose types are from the set Π . The arguments of the predicates of Σ may be both positive and negative literals, but a single variable cannot be contained in two distinct arguments of the same predicate. If a single variable is allowed to be contained in several distinct arguments of a single predicate, it is possible to enhance the set Π by Boolean predicates obtained from the predicates of Π by identifying some of their arguments. The goal is to find a truth assignment which satisfies the largest fraction $\rho(\Sigma)$ of the predicates of Σ . Hence, $\rho_l(\Pi) = \inf \rho(\Sigma)$ where the infimum is taken over all l -consistent sets Σ of (unweighted) predicates whose types are from the set Π . Similarly, if Σ is a set of weighted predicates, $\rho(\Sigma)$ denotes the ratio between the weights of the predicates which can be satisfied and the total weight of all the predicates of Σ and $\rho_l^w(\Pi) = \inf \rho(\Sigma)$ where the infimum is taken over all l -consistent sets Σ of weighted predicates. Note that in the unweighted case, Σ is a set, not a multiset (otherwise, the ratios ρ_∞ and ρ_∞^w would coincide).

If P is a Boolean predicate, $\sigma(P)$ denotes the number of combinations of arguments which satisfy P . Two Boolean predicates P and P' are *isomorphic* if they differ only by a permutation of the arguments and negations of some of them, e.g., if $P(x_1, x_2) = P'(x_2, \neg x_1)$, then the predicates P and P' are isomorphic. Clearly, if P and P' are two isomorphic predicates, then $\rho_l(P) = \rho_l(P')$ for all $l \geq 1$. A k -ary predicate is *essentially k -ary* if it depends on all its k arguments. If the predicate P is not essentially k -ary, it is isomorphic to a predicate P' such that $P'(x_1, \dots, x_k) = P''(x_1, \dots, x_{k-1})$ for some $(k-1)$ -ary Boolean predicate P'' . It is not hard to see that $\rho_l(P) = \rho_l(P') = \rho_l(P'')$ for all $l \geq 1$ in such case. Hence, in order to determine $\rho_l(P)$ for all unary, binary and ternary Boolean predicates P , it is enough to compute the values for representatives of isomorphism classes of essentially unary, binary and ternary Boolean predicates. The following three simple observations will be useful later:

Lemma 1 *Let P be a k -ary Boolean predicate P . It holds that $\rho_l(P) \geq \rho_l^w(P) \geq \sigma(P)/2^k$ for all $l \geq 1$.*

Proof: Let Σ be a set of N predicates of type P whose arguments are the variables x_1, \dots, x_n . Choose each of the variables x_i , $1 \leq i \leq n$, randomly and independently to be true with the probability $1/2$. Each predicate of the set Σ is satisfied by the constructed random truth assignment with probability $\sigma(P)/2^k$. Hence, the expected number of satisfied predicates is $N \cdot \sigma(P)/2^k$. Consequently, there is a truth assignment which satisfies at least $N \cdot \sigma(P)/2^k$ of Σ predicates and $\rho(\Sigma) \geq \sigma(P)/2^k$. □

Lemma 2 *It holds that $\rho_1(P) = \rho_1^w(P) = \sigma(P)/2^k$ for each k -ary Boolean predicate P .*

Proof: By Lemma 1, $\rho_1(P) \geq \rho_1^w(P) \geq \sigma(P)/2^k$. We construct a set Σ of predicates of type P with variables x_1, \dots, x_k and with $\rho(\Sigma) = \sigma(P)/2^k$. It is enough to set Σ to be the set consisting of all $P(x_1^{a_1}, \dots, x_k^{a_k})$ where $(a_1, \dots, a_k) \in \{0, 1\}^k$ and x_i^0 is $\neg x_i$ and x_i^1 is x_i . Clearly, each truth assignment satisfies exactly $\sigma(P)$ predicates out of all the 2^k predicates of Σ . Therefore, $\rho(\Sigma) = \sigma(P)/2^k$. □

Lemma 3 *If P be a k -ary predicate with $\sigma(P) = 1$, then $\rho_1(P) = \rho_1^w(P) = 2^{-k}$ and $\rho_l(P) = \rho_l^w(P) = 1$ for every $l \geq 2$.*

Proof: The equality $\rho_1(P) = 2^{-k}$ follows from Lemma 2. Let us consider a 2-consistent set Σ of predicates of type P . Since $\sigma(P) = 1$, each predicate of Σ forces the values to all its arguments. However, all the predicates must force the same value to a single variable because Σ is 2-consistent. Therefore, the “forced” truth assignment satisfies all the predicates of Σ and $\rho(\Sigma) = 1$. This immediately yields that $\rho_l(P) = 1$ for every $l \geq 2$. □

A *restriction* of a predicate P is a predicate P' obtained from P by fixing values of some of its arguments, e.g., $P'(x_1, x_2) = (x_1 \wedge x_2)$ is a restriction of the predicate $P(x_1, x_2, x_3) = (x_1 \wedge x_2 \wedge x_3) \vee (\neg x_3)$ obtained by fixing the value

of x_3 to be true. A restriction P' of a k -ary predicate P can be described by a vector $\tau \in \{0, 1, \star\}^k$ where 0 and 1 denote an argument which is fixed to be false and true, respectively, and \star denotes an unfixed argument. Let $\pi_{P,\tau}(p) : \langle 0, 1 \rangle \rightarrow \langle 0, 1 \rangle$ be equal to the probability that the k -ary predicate P with arguments x_1, \dots, x_k is satisfied if each x_i is set to be true randomly and independently with the probability $1 - p$, p and $1/2$, if τ_i is 0, 1 and \star , respectively. Note that $\pi_{P,\tau}(p)$ is a polynomial in p of degree at most k . For a set Π of predicates, let $\pi(\Pi)$ be the set of all the functions $\pi_{P,\tau}$ where $P \in \Pi$ and the restriction of P corresponding to τ is 1-extendable.

Example 1 Let Π be the set consisting of two predicates $P_1(x_1) = (x_1)$ and $P_2(x_1, x_2) = (x_1 \vee x_2)$. There is a single restriction of the predicate P_1 which is 1-extendable and this restriction corresponds to the vector 1. There are four restrictions of the predicate P_2 which are 1-extendable, those corresponding to 11, 1 \star , \star 1 and $\star\star$. Hence, the set $\pi(\Pi)$ consists of the following four functions:

$$\begin{aligned} \pi_{P_1,1}(p) &= p, & \pi_{P_2,11}(p) &= 2p - p^2, \\ \pi_{P_2,1\star}(p) &= \pi_{P_2,\star 1}(p) = (p + 1)/2, & \pi_{P_2,\star\star}(p) &= 3/4. \end{aligned}$$

Example 2 Consider a set Π containing the predicate $P(x_1, x_2, x_3, x_4, x_5) = (x_1 \wedge (x_2 \vee x_3 \vee x_4 \vee x_5))$. There are several restrictions of P which are 1-extendable, but each such restriction is isomorphic to a restriction corresponding to one of the following vectors: 1 $\star\star\star\star$, 10 $\star\star\star$, 11 $\star\star\star$, 100 $\star\star$, 110 $\star\star$, 111 $\star\star$, 1100 \star , 1100 \star , 1110 \star , 1111 \star , 11000, 11100, 11110 and 11111.

Let Ψ be the functional which assigns a continuous function $f : \langle 0, 1 \rangle \rightarrow \langle 0, 1 \rangle$ its maximum on the interval $\langle 0, 1 \rangle$. If F is a finite family of functions $f : \langle 0, 1 \rangle \rightarrow \langle 0, 1 \rangle$, then $\Psi(F)$ is defined to be the infimum $\Psi(f)$ where the function f ranges over all convex combinations of the functions of F . Note that the infimum is attained if the set F is a set of polynomials (which is the case of $\pi(\Pi)$ for any set of predicates Π). As mentioned in Section 1, one of our results is that the limit $\rho_\infty(\Pi) = \lim_{l \rightarrow \infty} \rho_l(\Pi)$ is equal to $\Psi(\pi(\Pi))$ for any set Π of Boolean predicates with arities at least two and $\rho_\infty^w(\Pi)$ is equal to $\Psi(\pi(\Pi))$ for any set Π of Boolean predicates (see Corollary 2 and Examples 3–6 after it). Let us draw the reader's attention to Example 6 which shows that it is not enough to consider only the function of $\pi(\Pi)$, but it is indeed necessary to consider their convex combinations.

3 Upper bound for the asymptotic case

Before we can design the algorithm for the asymptotic case, we first establish the following lemma on the derivatives of convex combinations of the functions contained in $\pi(\Pi)$:

Lemma 4 *Let Π be a set of predicates of arity at most K and let $f(p)$ be any convex combination of functions contained in $\pi(\Pi)$. The derivative of the function $f(p)$ for $p \in \langle 0, 1 \rangle$ takes values from the interval $\langle -K, +K \rangle$.*

Proof: Since the derivative of a convex combination of some functions is a convex combination of their derivatives, it is enough to prove the statement of the lemma only for the functions contained in the set $\pi(\Pi)$. Let f be a function contained in $\pi(\Pi)$ corresponding to a predicate $P \in \Pi$ and a vector τ . Let k be the arity of P (which is also the length of τ) and k' the number of 0's and 1's contained in τ . The function f can be expressed as the following linear combination:

$$f(p) = \sum_{i_1=0}^1 \cdots \sum_{i_{k'}=0}^1 \alpha_{i_1, \dots, i_{k'}} \prod_{j=1}^{k'} f_{i_j}(p)$$

where $0 \leq \alpha_{i_1, \dots, i_{k'}} \leq 1$, $f_0(p) = (1 - p)$ and $f_1(p) = p$. The derivative f' of f is the following:

$$\begin{aligned} f'(p) &= \sum_{i_1=0}^1 \cdots \sum_{i_{k'}=0}^1 \alpha_{i_1, \dots, i_{k'}} \sum_{j_0=1}^{k'} (-1)^{1+i_{j_0}} \prod_{j=1, j \neq j_0}^{k'} f_{i_j}(p) \\ &= \sum_{j_0=1}^{k'} \sum_{i_{j_0}=0}^1 (-1)^{1+i_{j_0}} \sum_{i_1=0}^1 \cdots \sum_{i_{j_0-1}=0}^1 \sum_{i_{j_0+1}=0}^1 \cdots \sum_{i_{k'}=0}^1 \alpha_{i_1, \dots, i_{k'}} \prod_{j=1, j \neq j_0}^{k'} f_{i_j}(p). \end{aligned}$$

It remains to estimate the absolute value of $f'(p)$ for $p \in \langle 0, 1 \rangle$:

$$\begin{aligned} |f'(p)| &\leq \sum_{j_0=1}^{k'} \left| \sum_{i_{j_0}=0}^1 (-1)^{1+i_{j_0}} \sum_{i_1=0}^1 \cdots \sum_{i_{j_0-1}=0}^1 \sum_{i_{j_0+1}=0}^1 \cdots \sum_{i_{k'}=0}^1 \alpha_{i_1, \dots, i_{k'}} \prod_{j=1, j \neq j_0}^{k'} f_{i_j}(p) \right| \\ &\leq \sum_{j_0=1}^{k'} 1 = k' \leq K. \end{aligned}$$

In order to establish the middle inequality, observe first that

$$\sum_{i_1=0}^1 \cdots \sum_{i_{j_0-1}=0}^1 \sum_{i_{j_0+1}=0}^1 \cdots \sum_{i_{k'}=0}^1 \prod_{j=1, j \neq j_0}^{k'} f_{i_j}(p) = 1$$

for all $p \in \langle 0, 1 \rangle$ and $j_0 = 1, \dots, k'$. Since both the function f_0 and f_1 are non-negative, the value of the function

$$\sum_{i_1=0}^1 \cdots \sum_{i_{j_0-1}=0}^1 \sum_{i_{j_0+1}=0}^1 \cdots \sum_{i_{k'}=0}^1 \alpha_{i_1, \dots, i_{k'}} \prod_{j=1, j \neq j_0}^{k'} f_{i_j}(p)$$

is always between 0 and 1 for $p \in \langle 0, 1 \rangle$, $j_0 = 1, \dots, k'$ and $i_{j_0} = 0, 1$. Since the absolute value of the difference of two numbers between 0 and 1 does not exceed 1, the inequality follows. □

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

Theorem 1 *Let Π be a fixed set of Boolean predicates and let K be the maximum arity of a predicate contained in Π . There exists an algorithm which given $\varepsilon > 0$ and a set of weighted predicates Σ of total weight w_0 either finds a truth assignment which satisfies predicates of Σ whose weight is at least $(\Psi(\pi(\Pi)) - \varepsilon)w_0$ or finds a set of at most $2K^{\lceil 2K/\varepsilon \rceil - 1}$ inconsistent predicates. Moreover, the algorithm runs in time linear in $|\Sigma|$ and $1/\varepsilon$.*

Proof: The algorithm consists of three steps:

1. Labeling variables according to the depth of “forcing” their values by the input predicates (or finding an inconsistent set which contains at most $2K^{\lceil 2K/\varepsilon \rceil - 1}$ predicates).
2. Finding a probability distribution on truth assignments such that the expected weight of the satisfied predicates is at least $(\Psi(\pi(\Pi)) - \varepsilon)w_0$.
3. Construction of a truth assignment which satisfies predicates whose weight is at least $(\Psi(\pi(\Pi)) - \varepsilon)w_0$.

The third step is an easy application of a standard linear-time derandomization technique proposed by Yannakakis [19] for locally consistent formulas (see also [11]) nowadays known as the method of conditional expectations (the reader is referred to [1, 2, 14] for additional details). So, we focus on the first two steps of the algorithm in the rest of the proof.

In the first step, we construct a sequence of $1 + \lceil 2K/\varepsilon \rceil$ partial truth assignments $\mu_0, \dots, \mu_{\lceil 2K/\varepsilon \rceil}$ and subsets $\Sigma_1, \dots, \Sigma_{\lceil 2K/\varepsilon \rceil}$ of Σ . The partial truth assignment μ_0 is the empty one, i.e., it sets no variables. Let i be an integer between 1 and $\lceil 2K/\varepsilon \rceil$ and assume that the partial truth assignment μ_0, \dots, μ_{i-1} have been constructed. Let Σ_i be the set of all the predicates of Σ whose restrictions with respect to μ_{i-1} are not 1-extendable. If there is a predicate whose restriction with respect to μ_{i-1} is constantly false, we stop. Otherwise, the partial truth assignment μ_{i-1} is extended to the partial truth assignment μ_i by setting the values of the variables forced by the restrictions of the predicates contained in Σ_i . The value of a variable x is *forced* if there exists a predicate which can be satisfied only if either x is false or x is true. If the value of a single variable is forced to be both true and false, we also stop.

Let us say few comments on the actual implementation of the first step of the algorithm. Each variable x will be labeled by the smallest i such that μ_i assigns the value to x . The variables whose values are forced by previously fixed variables are stored in a FIFO queue. When a variable is dequeued, the algorithm checks whether there are some new variables forced after fixing the value of the dequeued variable. If so, the newly forced variables are added to the end of the queue. In addition, in order to be able to quickly find inconsistent sets of clauses, we store for each variable which of the predicates forced its value and include this predicate to the corresponding set Σ_i . Note that the labels of the variables correspond to “depths” of derivations forcing their values and that each predicate is included to at most K of the sets $\Sigma_1, \dots, \Sigma_{\lceil 2K/\varepsilon \rceil}$.

If we stop because we find an unsatisfied predicate or a variable which is forced to two different values, we can easily construct an inconsistent set of at most $2(K^{\lceil 2K/\varepsilon \rceil - 1} + 1)$ predicates as described in the following. If an unsatisfied predicate is found, consider a set A consisting of this predicate, all the (at most K) predicates forcing the values of the variables contained in its arguments, all the (at most $K(K-1)$) predicates forcing the values of the variables contained in the “first-level” predicates, etc. Since there are at most $\lceil 2K/\varepsilon \rceil$ levels, the number of the predicates included to the set A does

not exceed:

$$1 + K + K(K - 1) + \dots + K(K - 1)^{\lceil 2K/\varepsilon \rceil - 2} \leq K^{\lceil 2K/\varepsilon \rceil - 1} + 1 .$$

If we stop because there is a variable which is forced to two different values, we include to the set A the two predicates which force it to have opposite values, all the (at most $2(K - 1)$) predicates forcing the values of the variables contained in their arguments, etc. The number of the predicates included to the set A does not exceed in this case:

$$2 + 2(K - 1) + 2(K - 1)^2 + \dots + 2(K - 1)^{\lceil 2K/\varepsilon \rceil - 2} \leq 2K^{\lceil 2K/\varepsilon \rceil - 1} .$$

In either of the cases, the number of the predicates contained in the set A is at most $2K^{\lceil 2K/\varepsilon \rceil - 1}$ and the set A can be constructed in time linear in $|A|K \leq |\Sigma|K$.

If for each variable x , a list of predicates which contain x is formed at the beginning of the computation (which can be simultaneously done for all the variables in linear time), the entire first step of the algorithm can be performed in time $O(|\Sigma|K)$ including the construction of an inconsistent set. Let us recall at this point that K is a constant since the set Π is fixed.

We now focus on the second step of the algorithm. Since each predicate of Σ can be contained in at most K sets $\Sigma_1, \dots, \Sigma_{\lceil 2K/\varepsilon \rceil}$, the total weight of all the predicates contained in the sets $\Sigma_1, \dots, \Sigma_{\lceil 2K/\varepsilon \rceil}$ when counting multiplicities does not exceed Kw_0 . By an averaging argument, there exists $1 \leq i \leq \lceil 2K/\varepsilon \rceil$ for which the weight of the predicates of Σ_i is at most $\varepsilon w_0/2$. Let w'_0 be the total weight of the predicates contained in $\Sigma \setminus \Sigma_i$. Note that $w'_0 \geq (1 - \varepsilon/2)w_0$ by the choice of i .

Let $f(p)$ be the expected weight of the satisfied predicates of $\Sigma \setminus \Sigma_i$ divided by w'_0 where each of the variables fixed by μ_{i-1} gets the value assigned to it by μ_{i-1} with the probability p and the remaining variables are set to be true with the probability $1/2$ (the values of all the variables are set mutually independently). Clearly, the coefficients of the polynomial $f(p)$ (of degree at most K) can be computed in time linear in $|\Sigma|$. Since the restriction of each predicate of $\Sigma \setminus \Sigma_i$ with respect to μ_{i-1} is 1-extendable, the function $f(p)$ is a convex combination of the functions from $\pi(\Pi)$. In particular, the absolute value of the derivative of $f(p)$ does not exceed K by Lemma 4.

Compute the value of the function $f(p)$ for each of the following values of p : $0, \frac{\varepsilon}{K}, \frac{2\varepsilon}{K}, \dots, \lfloor \frac{K}{\varepsilon} \rfloor \frac{\varepsilon}{K}, 1$. Let p_0 be the value for which the maximum is attained. Note that $f(p_0)$ differs from the maximum of the function $f(p)$ for

$p \in \langle 0, 1 \rangle$ by at most $\varepsilon/2$ because the absolute value of the derivative of f does not exceed K for $p \in \langle 0, 1 \rangle$. Since for each of the $\lfloor K/\varepsilon \rfloor + 2$ values of p , the function $f(p)$ can be evaluated in time $O(K)$, the algorithm needs time linear in $O(1/\varepsilon)$ to determine p_0 .

We claim that the probability distribution which assigns each of the variables fixed by μ_{i-1} the value assigned by μ_{i-1} with the probability p_0 and the remaining variables are set to be true with the probability $1/2$ is the desired probability distribution. The expected weight of the satisfied clauses is clearly at least $f(p_0)w'_0$. We further estimate this quantity:

$$f(p_0)w'_0 \geq \left(\max_{p \in \langle 0, 1 \rangle} f(p) - \varepsilon/2 \right) (1 - \varepsilon/2)w_0 \geq$$

$$(\Psi(\pi(\Pi)) - \varepsilon/2)(1 - \varepsilon/2)w_0 \geq (\Psi(\pi(\Pi)) - \varepsilon)w_0.$$

This finishes the second step of the algorithm. Let us point out that the algorithm does not need to compute any estimate on $\Psi(\pi(\Pi))$ in order to run correctly. □

An immediate corollary of Theorem 1 is the following:

Corollary 1 *Let Π be a set of Boolean predicates. For each $\varepsilon > 0$, there exists an integer $l \geq 1$ such that*

$$\rho_l(\Pi) \geq \rho_l^w(\Pi) \geq \Psi(\pi(\Pi)) - \varepsilon.$$

4 General lower bound

First, we introduce several concepts which are used throughout this section. If Σ is a set of predicates and μ is a partial truth assignment, then the *restriction* of Σ with respect to μ is the set Σ' of the predicates obtained from Σ by fixing the values of variables set by μ . The *dependence graph* $G(\Sigma')$ of a Σ' is the multigraph whose vertices are predicates of Σ' and the number of edges between two predicates P_1 and P_2 of Σ' is equal to the number of variables which appear in arguments of both the predicates P_1 and P_2 (regardless whether they appear as positive or negative literals). Note that the predicates whose arguments contain only the variables fixed by μ are isolated vertices in $G(\Sigma')$. A *semicycle* of length l of Σ with respect to μ is

a set Γ of l predicates such that the vertices corresponding to the predicates of Γ form a cycle of length l in $G(\Sigma')$. The following lemma relates the girth of the graph $G(\Sigma')$ and the local consistency of Σ for a suitable partial truth assignment μ :

Lemma 5 *Let Σ be a set of predicates, μ a partial truth assignment, Σ' the restriction of Σ with respect to μ and $l \geq 2$ an integer. If each predicate of Σ' is 1-extendable and Σ contains no semicycle of length at most l with respect to μ , then the set Σ is l -consistent.*

Proof: We prove by induction on i that any $i = 1, \dots, l$ predicates of Σ' can be simultaneously satisfied. This clearly implies the statement of the lemma because a truth assignment for Σ' can be viewed as an extension of the truth assignment μ to Σ .

The claim trivially holds for $i = 1$. Assume now that $i > 1$ and let P_1, \dots, P_i be any i predicates of Σ . Since $G(\Sigma')$ contains no cycle of length at most l , the vertices corresponding to P_1, \dots, P_i induce a forest T in $G(\Sigma')$. We can assume without loss of generality that P_i corresponds to a leaf or an isolated vertex in the forest T . Let y_1, \dots, y_n be the variables contained in the first $i - 1$ predicates which are not set by μ . By the induction hypothesis, there is a truth assignment for the variables y_1, \dots, y_n which satisfies all the predicates P_1, \dots, P_{i-1} . Since P_i is a leaf or an isolated vertex in T , it has at most one variable in common with the predicates P_1, \dots, P_{i-1} . Hence, the truth assignment for y_1, \dots, y_n can be extended to a truth assignment which satisfies all the predicates P_1, \dots, P_i because the restriction of the predicate P_i with respect to μ is 1-extendable. □

In the proof of the lower bound, Markov's inequality and Chernoff's inequality are used to bound the probability of large deviations from the expected value. The reader is referred to [7] for a more detailed exposition:

Proposition 1 *Let X be a non-negative random variable with the expected value E . The following holds for every $\alpha \geq 1$:*

$$\text{Prob}(X \geq \alpha) \leq \frac{E}{\alpha} .$$

Proposition 2 *Let X be a random variable equal to the sum of N zero-one independent random variables such that each of them is equal to 1 with the probability p . Then, the following holds for every $0 < \delta \leq 1$:*

$$\text{Prob}(X \geq (1 + \delta)pN) \leq e^{-\frac{\delta^2 pN}{3}} \quad \text{and} \quad \text{Prob}(X \leq (1 - \delta)pN) \leq e^{-\frac{\delta^2 pN}{2}}.$$

We are now ready to prove our lower bounds on $\rho_\infty^w(\Pi)$ and $\rho_\infty(\Pi)$:

Theorem 2 *Let Π be a set of Boolean predicates. For any integer $l \geq 1$ and any real $\varepsilon > 0$, there exists an l -consistent set Σ_0 of weighted predicates whose types are from the set Π such that:*

$$\rho^w(\Sigma_0) \leq \Psi(\pi(\Pi)) + \varepsilon.$$

Moreover, if the arity of each predicate Π is at least two, then there exists such a set Σ_0 of unweighted predicates.

Proof: We assume without loss of generality that $\varepsilon < 1$ is the inverse of a power of two. Let f_1, \dots, f_K be all the different functions contained in the set $\pi(\Pi)$ and let $\sum_{i=1}^K \alpha_i f_i$ be their convex combination with $\Psi(\sum_{i=1}^K \alpha_i f_i) = \Psi(\pi(\Pi))$. Let further P^i be a predicate of Π whose restriction with respect to a vector τ^i is 1-extendable and $\pi_{P^i, \tau^i} = f_i$. Observe that there are no two indices $i \neq i'$ such that $P^i = P^{i'}$ and $\tau^i = \tau^{i'}$. Finally, let K_0 be the maximum arity of a predicate contained in Π .

We consider a random set Σ of predicates whose arguments contain variables x_1, \dots, x_n and y_1, \dots, y_n where n is a sufficiently large power of two which will be fixed later in the proof. Fix an integer $i = 1, \dots, K$ and let k be the arity of P^i and k' the number of stars contained in τ^i . At this point, we abandon the condition that each variable can appear in at most one of the arguments of the predicate and we allow to include to Σ predicates which do not satisfy this condition. Later, we prune the set Σ to obey this constraint.

If $k > 1$, each of the $n^k 2^{k'}$ instances of the predicate P^i whose j -th argument, $1 \leq j \leq k$, is a positive literal containing one of the variables x_1, \dots, x_n if $\tau_j^i = 1$, a negative literal containing one of the variables x_1, \dots, x_n if $\tau_j^i = 0$ and a positive or negative literal containing one of the variables y_1, \dots, y_n if $\tau_j^i = \star$, is included to Σ randomly and independently of the other predicates with the probability $\alpha_i 2^{-k'} n^{-(k-1)+1/2l}$. The weights of all these predicates are set to one.

If $k = 1$, each instance of the predicate P^i whose only argument is a positive literal containing one of the variables x_1, \dots, x_n if $\tau_1^i = 1$, a negative literal containing one of the variables x_1, \dots, x_n if $\tau_1^i = 0$ and a positive or negative literal containing one of the variables y_1, \dots, y_n if $\tau_1^i = \star$, is included to Σ with the weight $\alpha_i 2^{-k'} n^{1/2l}$. Note that if the arity of each predicate of Π is at least two, the obtained set Σ consists of unweighted predicates (more precisely, all its predicates have the weight equal to one).

Let Σ^i be the predicates of Σ corresponding to P^i and τ^i . We prove the following three statements (under the assumption that n is sufficiently large):

1. The total weight of the predicates of Σ^i is at least $\alpha_i(1 - \frac{\varepsilon}{8})n^{1+1/2l}$ with the probability greater than $1 - 1/4K$.
2. With the probability greater than $1 - 1/4K$, each truth assignment which assigns true to exactly n' of the variables x_1, \dots, x_n satisfies the predicates of Σ^i whose total weight is at most $\alpha_i(f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l}$.
3. The total weight of the predicates whose arguments do not contain different variables is at most $\alpha_i \frac{\varepsilon}{8} n^{1+1/2l}$ with the probability greater than $1 - 1/4K$.

If the arity k of P^i is one or $\alpha_i = 0$, then all the three statements hold with the probability one. In the rest, we consider the case that the arity of P^i is at least two, i.e., $k \geq 2$, and $\alpha_i > 0$.

The probability that the total weight of the predicates of Σ^i is smaller than $\alpha_i(1 - \frac{\varepsilon}{8})n^{1+1/2l}$ is bounded by Proposition 2 from above by the following:

$$e^{-\frac{(\varepsilon/4)^2 (\alpha_i 2^{-k'} n^{-(k-1)+1/2l}) (n^k 2^{k'})}{2}} = e^{-\frac{\varepsilon^2 \alpha_i n^{1+1/2l}}{128}}.$$

Since ε , α_i , l and K do not depend on n , the probability that the total weight of the predicates of Σ^i exceeds $\alpha_i(1 - \frac{\varepsilon}{8})n^{1+1/2l}$ is smaller than $1/4K$ if n is sufficiently large.

Let μ be any of the 2^{2n} truth assignments for the variables x_1, \dots, x_n and y_1, \dots, y_n ; let n' be the number of variables x_1, \dots, x_n which are set to be true by μ . A predicate which can be included to Σ^i is said to be *good* if it is satisfied by μ . Note that there are exactly $f_i(n'/n)n^k 2^{k'}$ good predicates. If $f_i(n'/n) \leq \frac{\varepsilon}{8}$, then mark additional predicates to be good so that the total number of good predicates is $\frac{\varepsilon}{8} n^k 2^{k'}$ (note that since ε is the inverse of a power of two, then this expression is an integer if n is a sufficiently large

power of two). Hence, the expected number of good predicates included to Σ^i is exactly $\max\{f_i(n'/n), \varepsilon/8\}n^k2^{k'} \cdot \alpha_i n^{-(k-1)+1/2l}2^{-k'}$. Using the fact that $f_i(n'/n) \leq 1$ and Proposition 2, we infer the following:

$$\begin{aligned} & \text{Prob}(\mu \text{ satisfies more than } \alpha_i(f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l} \text{ predicates of } \Sigma^i) \leq \\ & \text{Prob}(\Sigma^i \text{ contains more than } \alpha_i(f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l} \text{ good predicates)} \leq \\ & \text{Prob}(\Sigma^i \text{ contains } > (1+\varepsilon/8)\alpha_i \max\{f_i(n'/n), \varepsilon/8\}n^{1+1/2l} \text{ good predicates)} \leq \\ & e^{-\frac{\varepsilon^2 \alpha_i \max\{f_i(n'/n), \varepsilon/8\}n^{1+1/2l}}{192}} \leq e^{-\frac{\varepsilon^3 \alpha_i n^{1+1/2l}}{1536}}. \end{aligned}$$

Since there are 2^{2n} possible truth assignment μ , the probability that there exists one which satisfies more than $\alpha_i(f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l}$ clauses of Σ^i is at most $2^{2n} \cdot e^{-\frac{\varepsilon^3 \alpha_i n^{1+1/2l}}{1536}}$. Since ε , α_i and K are fixed, this probability is smaller than $1/4K$ if n is sufficiently large.

It remains to establish our third claim on Σ^i . At most $\binom{k}{2}n^{k-1}2^{k'}$ out of all the $n^k2^{k'}$ predicates which can be included to Σ^i contain one variable in several of its arguments. Therefore, the expected number of such predicates which are contained in the set Σ^i is at most $\binom{k}{2}n^{k-1}2^{k'}\alpha_i2^{-k'}n^{-(k-1)+1/2l} = \alpha_i\binom{k}{2}n^{1/2l}$. By Markov's inequality (Proposition 1), the probability that the number of such predicates in Σ^i exceeds $\alpha_i\frac{\varepsilon}{8}n^{1+1/2l}$ is at most the following fraction:

$$\frac{\alpha_i\binom{k}{2}n^{1/2l}}{\alpha_i\frac{\varepsilon}{8}n^{1+1/2l}} = \binom{k}{2}\frac{8}{\varepsilon n}.$$

Since ε , k and K are independent of n , the probability of this event is smaller than $1/4K$ if n is sufficiently large.

It can be concluded that with the probability greater than $1/4$ the following three statements hold for the set Σ and a sufficiently large n (recall that $\sum_{i=1}^K \alpha_i = 1$):

1. The total weight of the predicates of Σ is at least $(1 - \frac{\varepsilon}{8})n^{1+1/2l}$.
2. Any truth assignment which assigns true to exactly n' of the variables x_1, \dots, x_n satisfies the predicates of Σ whose total weight does not exceed $(\sum_{i=1}^K \alpha_i f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l}$.
3. The total weight of the predicates whose arguments do not contain different variables is at most $\frac{\varepsilon}{8}n^{1+1/2l}$.

We now estimate the number of semicycles of length at most l in Σ with respect to the partial truth assignment μ_0 which sets all the variables x_1, \dots, x_n to be true. Note that all the restrictions of the predicates contained in Σ with respect to μ_0 are 1-extendable. Let us consider a semicycle corresponding to the predicates $P'_1, \dots, P'_{l'}$, $2 \leq l' \leq l$, described by $\tau'_1, \dots, \tau'_{l'}$. Let k_i be the arity of the predicate P'_i and k'_i the number of stars in τ'_i . The number of all semicycles corresponding to the restrictions of the predicates $P'_1, \dots, P'_{l'}$ determined by $\tau'_1, \dots, \tau'_{l'}$ is at most $\prod_{i=1}^{l'} n^{k_i - k'_i} n^{k'_i - 1} 2^{k'_i} k'_{i-1}$ (the indices are taken modulo l' , i.e., $k'_0 = k'_{l'}$). The probability of including any such particular sequence to Σ is $\prod_{i=1}^{l'} \alpha'_i n^{-(k_i + k'_i - 1) + 1/2l} 2^{-k'_i}$ where α'_i is the coefficient α_i corresponding to P'_i and τ'_i . Therefore, the expected number of semicycles contained in Σ which correspond to the restrictions of the predicates $P'_1, \dots, P'_{l'}$ determined by $\tau'_1, \dots, \tau'_{l'}$ is at most $\prod_{i=1}^{l'} k'_i n^{1/2l} \leq K'_0 n^{1/2}$ (recall that $0 \leq \alpha'_i \leq 1$ for all $1 \leq i \leq l'$ and K_0 denotes the maximum arity of a predicate in Π).

Since there are at most $K^{l'}$ ways how to choose the predicates $P'_1, \dots, P'_{l'}$ and $3^{K_0 l'}$ possible choices of the vectors $\tau'_1, \dots, \tau'_{l'}$, the expected number of semicycles of Σ of length l' does not exceed $(KK_0 3^{K_0})^{l'} n^{1/2}$. By Proposition 1, the probability that Σ contains more than $\frac{\varepsilon}{8l} n^{1+1/2l}$ semicycles of length at most l is at most the following:

$$\frac{l(KK_0 3^{K_0})^l n^{1/2}}{\frac{\varepsilon}{8l} n^{1+1/2l}} \leq \frac{8l^2 (KK_0 3^{K_0})^l}{\varepsilon n^{1/2}}.$$

Since the numbers l , K , K_0 and ε do not depend on n , this probability is smaller than $1/4$ if n is sufficiently large. Therefore with positive probability, the set Σ has the properties 1–3 stated above and the number of its semicycles of length at most l with respect to the partial truth assignment μ_0 is at most $\frac{\varepsilon}{8l} n^{1+1/2l}$. For the rest of the proof, fix Σ' to be any such set of predicates.

Remove from the set Σ' all the predicates contained in semicycles of length at most l with respect to μ_0 and all the predicates which contains the same variable in several of their arguments. Let Σ_0 be the resulting set of predicates. Note that there are at most at most $l \cdot \frac{\varepsilon}{8l} n^{1+1/2l} = \frac{\varepsilon}{8} n^{1+1/2l}$ predicates contained in semicycles of length at most l . Since each of the predicates of Σ' which is contained in a semicycle must contain one of the variables y_1, \dots, y_n , its arity is at least two. Consequently, its weight is equal to one. Hence, the total weight of the predicates removed from Σ' is at most $\frac{\varepsilon}{8} n^{1+1/2l} + \frac{\varepsilon}{8} n^{1+1/2l} = \frac{\varepsilon}{4} n^{1+1/2l}$ and the total weight of the predicates of Σ_0

is at least $(1 - \frac{3\varepsilon}{8})n^{1+1/2l}$. Clearly, the total weight of the predicates of Σ_0 which can be simultaneously satisfied by a truth assignment does not exceed the total weight of such predicates of Σ' . We can now conclude that the following holds for each truth assignment which sets n' ($0 \leq n' \leq n$) of the variables x_1, \dots, x_n to be true:

$$\rho^w(\Sigma_0) \leq \frac{(\sum_{i=1}^K \alpha_i f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l}}{(1 - \frac{3\varepsilon}{8})n^{1+1/2l}} \leq \frac{\Psi(\pi(\Pi)) + \frac{\varepsilon}{4}}{1 - \frac{3\varepsilon}{8}} \leq$$

$$\Psi(\pi(\Pi)) \frac{1 + \frac{\varepsilon}{4}}{1 - \frac{3\varepsilon}{8}} \leq \Psi(\pi(\Pi))(1 + \varepsilon) \leq \Psi(\pi(\Pi)) + \varepsilon.$$

Since Σ_0 contains no semicycles of length at most l with respect to μ_0 and all the restrictions of the predicates of Σ_0 with respect to μ_0 are 1-extendable, the set Σ_0 is l -consistent by Lemma 5. Consequently, $\rho_l^w(\Pi) \leq \Psi(\pi(\Pi)) + \varepsilon$. Moreover, if the arity of each predicate of Π is at least two, the weights of all the predicates of Σ are one and $\rho_l(\Pi) \leq \Psi(\pi(\Pi)) + \varepsilon$. □

We immediately infer from Corollary 1 and Theorem 2 the following expressions for $\rho_\infty(\Pi)$ and $\rho_\infty^w(\Pi)$:

Corollary 2 *Let Π be a finite set of Boolean predicates. The following holds:*

$$\rho_\infty^w(\Pi) = \Psi(\pi(\Pi)).$$

Moreover, if the arity of each predicate of Π is at least two, then the following holds:

$$\rho_\infty(\Pi) = \Psi(\pi(\Pi)).$$

As an application of Corollary 2, we compute the values $\rho_\infty^w(\Pi)$ for several sets Π :

Example 3 *Let Π be the set of predicates from Example 1. Since $\pi_{P_2, \star\star}(p)$ equals to $3/4$ for all $0 \leq p \leq 1$, we infer $\Psi(\pi(\Pi)) \leq \Psi(\pi_{P_2, \star\star}) = 3/4$. On the other hand, the value of each of the functions $\pi_{P_1, 1}$, $\pi_{P_2, 11}$, $\pi_{P_2, 1\star}$ and $\pi_{P_2, \star\star}$ for $p = 3/4$ is at least $3/4$. Thus, the value of any convex combination of them for $p = 3/4$ is also at least $3/4$ and $\Psi(\pi(\Pi)) \geq 3/4$. Hence, $\rho_\infty^w(\Pi) = 3/4$.*

Example 4 Let Π be the set of predicates from Example 2. Since the function $\pi_{P,100\star\star}(p)$ is $p(1-p^2/4)$, we infer that $\Psi(\pi(\Pi)) \leq \Psi(\pi_{P,100\star\star}) = 3/4$. On the other hand, each of the functions $\pi_{P,\tau}$ for all the vectors τ from Example 2 is at least $3/4$ for $p = 1$. Therefore, $\rho_\infty^w(\Pi) = \Psi(\pi(\Pi)) = 3/4$.

Example 5 Let Π^k be the set containing a single predicate $P(x_1, \dots, x_k) = x_1 \wedge (x_2 \vee \dots \vee x_k)$ for an integer $k \geq 7$. Consider the vector $\tau = 10 \dots 0 \star \star$. Clearly, the restriction of P determined by τ is 1-extendable. It is easy to show that the maximum of the function $\pi_{P,\tau}$ is attained for $p_0 = \sqrt[k-3]{\frac{4}{k-2}}$ and it is strictly larger than $3/4$. Moreover, the value $\pi_{P,\tau}(p_0)$ is smaller or equal to the value $\pi_{P,\tau'}(p_0)$ for any τ' corresponding to a 1-extendable restriction of P . We infer that $\rho_\infty^w(\Pi^k) = \Psi(\pi(\Pi^k)) \geq \Psi(\pi_{P,\tau}) > 3/4$.

In our final example, we show that it is necessary to consider convex combinations of the polynomials of $\pi(\Pi)$ in the expression for $\rho_\infty^w(\Pi)$, i.e., we present an example of a set Π where the value of the functional Ψ for each function from $\pi(\Pi)$ is strictly larger than $\Psi(\pi(\Pi))$.

Example 6 Let Π be a set consisting of two predicates $P_1(x_1) = x_1$ and $P_2(x_1, x_2, x_3) = x_1 \vee x_2 \vee x_3$. It is straightforward to verify that $\Psi(f) \geq 7/8$ for each polynomial $f \in \pi(\Pi)$. Consider the following polynomial f_0 :

$$f_0(p) = \frac{1}{5}\pi_{P_1,1}(p) + \frac{4}{5}\pi_{P_2,0\star\star}(p) = \frac{p}{5} + \frac{4}{5}\left(1 - \frac{p}{4}\right) = \frac{4}{5}.$$

Clearly, f_0 is a convex combination of the polynomials of $\pi(\Pi)$. Therefore, $\Psi(\pi(\Pi)) \leq 4/5$. On the other hand, the value of each polynomial of $\pi(\Pi)$ for $p = 4/5$ is at least $4/5$. Consequently, $\Psi(\pi(\Pi)) \geq 4/5$. Corollary 2 now yields that $\rho_\infty^w(\Pi) = 4/5$.

Using Theorem 2, we are able to determine the values $\rho_l(P)$ for every 1-extendable Boolean predicate P and for every $l \geq 1$:

Theorem 3 If P is a k -ary Boolean predicate which is 1-extendable, then $\rho_l(P) = \rho_l^w(P) = \sigma(P)/2^k$ for each $l \geq 1$.

Proof: If $l = 1$, the statement follows from Lemma 2. Assume that $l \geq 2$. By Lemma 1, $\rho_l(P) \geq \sigma(P)/2^k$. We show that $\Psi(\pi(\Pi)) = \sigma(P)/2^k$. Since the function $\pi_{P,\tau}(p)$ with $\tau = \star, \dots, \star$ is constantly equal to $\sigma(P)/2^k$, the value

of $\Psi(\pi(\Pi))$ does not exceed $\sigma(P)/2^k$. On the other hand, $p(1/2) = \sigma(P)/2^k$ for every polynomial $p \in \pi(\Pi)$. Therefore, $\Psi(\pi(\Pi)) = \sigma(P)/2^k$ and the inequality $\rho_l(P) \leq \sigma(P)/2^k$ now follows from Theorem 2 (note that the arity of P is at least two unless it is constantly true). □

5 2-CNF formulas

In this section, we study structure of 2-CNF formulas, i.e., CNF formulas of clauses of sizes one and two. We first recall a well-known lemma about unsatisfiable formulas with clauses of sizes two which can be found, e.g., in [3]. If Φ is a 2-CNF formula with variables x_1, \dots, x_n , then $G(\Phi)$ denotes the directed graph of order $2n$ whose vertices correspond to literals x_1, \dots, x_n and $\neg x_1, \dots, \neg x_n$ and whose edge set is the following: For each clause $(a \vee b)$, $G(\Phi)$ contains an arc from the literal $\neg a$ to the literal b and an arc from $\neg b$ to a (note that both a and b represent here literals, not variables). For each clause (a) (which can also be viewed as a clause $(a \vee a)$), we include an arc from the literal $\neg a$ to a .

Lemma 6 *Let Φ be a 2-CNF formula with variables x_1, \dots, x_n . The formula Φ is satisfiable if and only if $G(\Phi)$ contains no directed cycle through both the vertices x_i and $\neg x_i$ for any i , $1 \leq i \leq n$.*

An immediate corollary of Lemma 6 is the following:

Lemma 7 *Each minimal inconsistent set of clauses of a 2-CNF contains at most two clauses of size one.*

We now show that there exist extremal 2-CNF formulas in which each small inconsistent set of clauses contains two clauses of sizes one:

Lemma 8 *Let $2 \leq l \leq L$ be any two integers. For each $\varepsilon > 0$, there exists a 2-CNF l -consistent formula Φ with $\rho(\Phi) \leq \rho_l(\Pi_{2\text{-SAT}}) + \varepsilon$ such that each inconsistent set of L clauses contains at least two clauses of size one. Moreover, Φ contains each single clause of size two at most once.*

Proof: Fix integers $l \geq 2$ and $L \geq l$ for the rest of the proof. Similarly, $\delta < 1$ is a positive real which will be chosen at the end of the proof. Fix an l -consistent formula Φ_0 with $\rho(\Phi_0) < \rho_l(\Pi_{2\text{-SAT}})(1 + \delta)$. We now classify the variables contained in the formula Φ_0 : The set A_1 is formed by variables x contained in a clause of size one; we can assume without loss of generality that each variable $x \in A_1$ appears as a positive literal in the clause of size one. The set A_i , $2 \leq i \leq l/2$, consists of variables x which are not contained in any A_j , $1 \leq j \leq i - 1$, and which are contained in a clause of the form $(\neg y \vee x)$ for $y \in A_{i-1}$. Since Φ is l -consistent, we can assume that all the occurrence of $x \in A_i$ in the clauses $(\neg y \vee x)$, $y \in A_{i-1}$, are positive: Otherwise, there would be a set of at most i clauses which force x to be true as well as a set of at most i clauses which force x to be false. The union of these two sets of clauses consists of at most $2i$ clauses and it is clearly inconsistent. Since $i \leq l/2$, this is impossible. Finally, let A_0 be the set of the remaining variables of Φ .

Let w_{ij} , $w_{\bar{i}j}$ and $w_{\bar{i}\bar{j}}$ be the number (sum of the weights) of the clauses of the type $(x \vee y)$, $(\neg x \vee y)$ and $(\neg x \vee \neg y)$, respectively, where $x \in A_i$ and $y \in A_j$. Similarly, let w_1 be the number (sum of the weights) of the clauses of the type (x) where $x \in A_1$. We may assume that $w_1 > 0$. Otherwise, $\rho(\Phi_0) \geq 3/4$ and we can set Φ to be an L -consistent set Σ with $P(x, y) = (x \vee y)$ with $\rho(\Sigma) \leq 3/4 + \varepsilon$ constructed in Theorem 3. Finally, W denotes the sum of all w_1 , w_{ij} , $w_{\bar{i}j}$ and $w_{\bar{i}\bar{j}}$ for $0 \leq i, j \leq l/2$. By the definition of the sets $A_1, \dots, A_{\lfloor l/2 \rfloor}$, $w_{\bar{i}j} = 0$ for all $1 \leq i, j \leq l/2$ with $i + 1 < j$. In addition, since Φ is l -consistent, $w_{\bar{i}\bar{j}} = 0$ for all $1 \leq i, j \leq l/2$ with $i + j + 1 \leq l$.

We now define W_p to be the maximum of the sum:

$$\begin{aligned} w_1 p_1 &+ \sum_{0 \leq i < j \leq l/2} w_{ij} (p_i + p_j - p_i p_j) + w_{\bar{i}\bar{j}} (1 - p_i p_j) \\ &+ \sum_{0 \leq i, j \leq l/2} w_{\bar{i}j} (1 - p_i + p_i p_j) \end{aligned} \quad (1)$$

where the maximum is taken over all $0 \leq p_0, \dots, p_{\lfloor l/2 \rfloor} \leq 1$. Clearly, $W_p/W \leq \rho(\Phi_0)$: Consider the probabilities $p_0, \dots, p_{\lfloor l/2 \rfloor}$ for which the maximum in the above expression is attained. If each of the variables of the set A_i , $0 \leq i \leq l/2$, is chosen to be true randomly and independently with the probability p_i , then the expected number (weight) of the satisfied clauses is W_p . Therefore, there is a truth assignment which satisfies at least this number of clauses and consequently $W_p/W \leq \rho(\Phi_0)$.

Let n be an integer which we fix later. Let X_i , $0 \leq i \leq l/2$, be $\lfloor l/2 \rfloor + 1$ disjoint sets consisting of n variables each. We construct a 2-CNF formula Φ with variables $X_0 \cup \dots \cup X_{\lfloor l/2 \rfloor}$. The formula Φ contains $n^{1/2L}$ copies of a clause (x) for each $x \in X_1$. The other clauses are included to the formula Φ randomly and independently as follows: The clauses $(x \vee y)$, $(\neg x \vee y)$ and $(\neg x \vee \neg y)$ where $x \in X_i$ and $y \in X_j$ with $i \neq j$ are included to Φ with the probability $w_{ij}n^{-1+1/2L}/w_1$, $w_{\bar{i}j}n^{-1+1/2L}/w_1$ and $w_{\bar{i}\bar{j}}n^{-1+1/2L}/w_1$, respectively. The clauses $(x \vee y)$, $(\neg x \vee y)$ and $(\neg x \vee \neg y)$ where $x, y \in X_i$ are included to Φ with the probability $2w_{ij}n^{-1+1/2L}/w_1$, $w_{\bar{i}j}n^{-1+1/2L}/w_1$ and $2w_{\bar{i}\bar{j}}n^{-1+1/2L}/w_1$, respectively.

We claim that the number of clauses of Φ is at least $Wn^{1+1/2L}(1 - \delta)/w_1$ and the number of clauses which can be simultaneously satisfied does not exceed $W_p n^{1+1/2L}(1 + \delta)/w_1 + 3Wn^{1+1/2L}\delta/w_1$ with the probability which tends to 1 as n goes to infinity. We show that each of the complementary events, i.e., the “bad” events, for each separate types of clauses occurs with the probability which tends to 0. Since the number of bad events is finite (and independent of n), this yields the claim. As an example, we present the analysis only for a single type of clauses, e.g., clauses $(x \vee y)$ for $x \in X_i$ and $y \in X_j$ with $i \neq j$ for fixed integers i and j . Namely, we aim to show that the number of clauses of this type is smaller than $w_{ij}n^{1+1/2L}(1 - \delta)/w_1$ with probability tending to 0. In addition, the probability that there is a truth assignment which assigns the true value to a fraction p_i, p_j , of the variables of X_i, X_j , respectively, and which satisfies more than $w_{ij}n^{1+1/2L}(1 - (1 - p_i)(1 - p_j))(1 + \delta)/w_1 + 3w_{ij}n^{1+1/2L}\delta/w_1 = w_{ij}n^{1+1/2L}(p_i + p_j - p_i p_j)(1 + \delta)/w_1 + 3w_{ij}n^{1+1/2L}\delta/w_1$ clauses of the considered type also tends to zero.

By Proposition 2, the probability that the number of clauses $(x \vee y)$ with $x \in X_i$ and $y \in X_j$ is smaller than $w_{ij}n^{1+1/2L}(1 - \delta)/w_1$ is at most:

$$e^{-\frac{\delta^2 w_{ij} n^{1+1/2L} / w_1}{3}} = e^{-\Theta(n^{1+1/2L})} \rightarrow 0.$$

The second part of the statement is more difficult. We first prove the claim for p_i and p_j where p_i or p_j is at least δ . Fix now a truth assignment for $X_i \cup X_j$ which assigns the true value to a fraction p_i, p_j , of the variables X_i, X_j , respectively. By Proposition 2, the probability that the number of clauses of the considered type satisfied by this fixed truth assignment exceeds $w_{ij}n^{1+1/2L}(p_i + p_j - p_i p_j)(1 + \delta)/w_1$ is at most:

$$e^{-\frac{\delta^2 \cdot w_{ij} n^{-1+1/2L} / w_1 \cdot n^2 (p_i + p_j - p_i p_j)}{3}} \leq e^{-\frac{\delta^3 \cdot w_{ij} n^{1+1/2L} / w_1}{3}} = e^{-\Theta(n^{1+1/2L})}.$$

Since there are at most 2^{2n} possible truth assignments the probability that there exists a truth assignment with $\delta \leq \max\{p_i, p_j\}$, which satisfy more clauses than claimed is at most $2^n e^{-\Theta(n^{1+1/2L})} \rightarrow 0$. We now show that if there is no “bad” truth assignment with $\delta \leq \max\{p_i, p_j\}$, then there is no “bad” truth assignment with $0 \leq p_i, p_j \leq \delta$. Consider a truth assignment which assigns the true value to at most δn variables of each X_i and X_j and modify it to a truth assignment which assigns the true value to $\lceil \delta n \rceil$ variables of each X_i and X_j . This modification can only increase the number of satisfied clauses of the considered type. Since both the modified p_i and p_j are now larger than δ , the assignment satisfies at most $w_{ij} n^{1+1/2L} (2\delta + 2/n - \delta^2)(1 + \delta)/w_1$ clauses. If n is sufficiently large, then this expression is at most $w_{ij} n^{1+1/2L} 3\delta/w_1$ as desired. This finishes the proof of the claim.

The number of variables of the formula Φ is at most $N = (\lfloor l/2 \rfloor + 1)n$. By Lemma 6, there are at most $(2N)^k$ minimal inconsistent sets of k clauses such that the size of each clause is two and there are at most $(2N)^{k-1}$ minimal inconsistent sets of $k-1$ clauses such that the size of each clause is two except for precisely one clause whose size is one. We omit the straightforward but little technical argument yielding these upper bounds due to space limitations. Since each clause of size two is included to Φ with the probability at most $W_p n^{-1+1/2L}/w_1$, the expected number of minimal inconsistent sets of at most L clauses containing zero or one clause of size one is at most the following:

$$\begin{aligned} & \sum_{k=1}^L \frac{(2N)^k W_p n^{-k+k/2L}}{w_1} + \frac{(2N)^{k-1} W_p n^{-(k-1)+(k-1)/2L}}{w_1} = \\ & \sum_{k=1}^L \frac{(2(\lfloor l/2 \rfloor + 1))^k W_p n^{k/2L}}{w_1} + \frac{(2(\lfloor l/2 \rfloor + 1))^{k-1} W_p n^{(k-1)/2L}}{w_1} \leq \\ & \sum_{k=1}^L \frac{(l+2)^k W_p n^{k/2L}}{w_1} + \frac{(l+w)^{k-1} W_p n^{(k-1)/2L}}{w_1} \leq \frac{2L(l+1)^L W_p n^{1/2}}{w_1}. \end{aligned}$$

By Markov’s inequality, the probability that there are more than $4L(l+1)^L W_p n^{1/2}/w_1$ such minimal inconsistent sets of clauses is at most $1/2$. Therefore, if n is sufficiently large (with respect to a previously fixed $\delta > 0$), with positive probability, the random formula Φ has at least $W n^{1+1/2L} (1 - \delta)/w_1$ clauses, at most $W_p n^{1+1/2L} (1 + \delta)/w_1 + 3W n^{1+1/2L} \delta/w_1$ clauses of Φ can be

simultaneously satisfied and Φ contains at most $4L(l+1)^L W_p n^{1/2}/w_1$ inconsistent sets of at most L clauses with no or a single clause of size one. Fix such a formula Φ . We obtain Φ' from Φ by removing all (at most) $4L^2(l+1)^L W_p n^{1/2}/w_1$ clauses of size two contained in an inconsistent set of at most L clauses with no or a single clause of size one. This may decrease the number of clauses of Φ by at most $4L^2(l+1)^L W_p n^{1/2}/w_1$. On the other hand, the number of clauses which can be simultaneously satisfied cannot increase.

We now estimate $\rho(\Phi')$ (observe that $W_p \geq W/2$):

$$\begin{aligned} \rho(\Phi') &\leq \frac{W_p n^{1+1/2L}(1+\delta)/w_1 + 3W n^{1+1/2L}\delta/w_1}{W n^{1+1/2L}(1-\delta)/w_1 - 4L^2(l+1)^L W_p n^{1/2}/w_1} = \\ &\frac{W_p(1+\delta) + 3\delta W}{W(1-\delta) - 4L^2(l+1)^L W_p n^{-1/2-1/2L}} \leq \frac{W_p(1+7\delta)}{W(1-\delta) - O(n^{-1/2-1/2L})}. \end{aligned}$$

Therefore, if n is sufficiently large, then:

$$\rho(\Phi') \leq \frac{W_p(1+7\delta)}{W(1-2\delta)} \leq \rho(\Phi_0) \frac{1+7\delta}{1-2\delta} \leq \rho_i(\Pi_{2\text{-SAT}}) \frac{(1+\delta)(1+7\delta)}{1-2\delta}.$$

Hence, for each $\varepsilon > 0$, we can choose $\delta > 0$ small enough that $\rho(\Phi') \leq \rho_i(\Pi_{2\text{-SAT}}) + \varepsilon$.

It remains to show that the formula Φ' is l -consistent. By Lemma 7, each minimal inconsistent set of clauses of Φ' contains at most two clauses of size one. On the other hand, each inconsistent set of at most L clauses contains at least two such clauses. Therefore, each minimal inconsistent set of at most l clauses of Φ' contains precisely two clauses of size one. Fix such a set Γ of clauses of Φ' and let (x_1) and (y_1) be the two clauses of size one contained in Γ . Obviously, $x_1, y_1 \in X_1$. By Lemma 6, Γ contains clauses of size two in which x_1 and y_1 appear as negative literals. By the construction of Φ' , such clauses can be only $(\neg x_1 \vee x_2)$ and $(\neg y_1 \vee y_2)$ for some $x_2, y_2 \in X_2$. By Lemma 6, Γ has to contain clauses of size two in which x_2 and y_2 appear as negative literals. By the construction of Φ' , such clauses can be only $(\neg x_2 \vee x_3)$ and $(\neg y_2 \vee y_3)$ for some $x_3, y_3 \in X_3$. In this way, we continue until we reach the set $X_{\lfloor l/2 \rfloor}$. By the minimality of the set Γ , $x_i \neq y_i$ for all $1 \leq i \leq \lfloor l/2 \rfloor$. Therefore, if $|\Gamma| \leq 2\lfloor l/2 \rfloor + 1$, then Γ contains the clauses $(x_1), (\neg x_1 \vee x_2), \dots, (\neg x_{\lfloor l/2 \rfloor - 1} \vee x_{\lfloor l/2 \rfloor}), (y_1), (\neg y_1 \vee y_2), \dots, (\neg y_{\lfloor l/2 \rfloor - 1} \vee y_{\lfloor l/2 \rfloor})$ and $(\neg x_{\lfloor l/2 \rfloor} \vee \neg y_{\lfloor l/2 \rfloor})$. Hence, $|\Gamma| > l$ if l is even. If l is odd, then $w_{\lfloor l/2 \rfloor \lfloor l/2 \rfloor} =$

0 and thus Φ' cannot contain the clause $(\neg x_{\lfloor l/2 \rfloor} \vee \neg y_{\lfloor l/2 \rfloor})$. This yields $|\Gamma| > l$ in this case, too. □

A close inspection of the proof of Lemma 8 yields that for any weights $w_1, w_{ij}, w_{\bar{i}j}$ and $w_{\bar{i}\bar{j}}$ with $w_{\bar{i}j} = 0$ for all $1 \leq i \leq j - 1$ and $w_{\bar{i}\bar{j}} = 0$ for all $1 \leq i, j$ with $i + j + 1 \leq l$. there is an l -consistent formula Φ with $\rho(\Phi) \leq W_p/W + \varepsilon$ where $W = w_1 + \sum_{i,j}(w_{ij} + w_{\bar{i}j} + w_{\bar{i}\bar{j}})$ and W_p is the maximum of the sum (1) taken over all $0 \leq p_0, \dots, p_{\lfloor l/2 \rfloor} \leq 1$. Therefore, we have the following formula for $\rho_l(\Pi_{2\text{-SAT}})$ for all $l \geq 2$:

Corollary 3 *For each $l \geq 2$, the following holds:*

$$\rho_l(\Pi_{2\text{-SAT}}) = \min_{\substack{0 \leq w_1, w_{ij}, w_{\bar{i}j}, w_{\bar{i}\bar{j}} \\ w_1 + \sum_{i,j}(w_{ij} + w_{\bar{i}j} + w_{\bar{i}\bar{j}}) = 1}} W_p$$

where the minimum is taken over all combinations of weights with $w_{\bar{i}j} = 0$ for all $1 \leq i \leq j - 1$ and $w_{\bar{i}\bar{j}} = 0$ for all $1 \leq i, j$ with $i + j + 1 \leq l$ and W_p is the maximum of the sum (1) taken over all $0 \leq p_0, \dots, p_{\lfloor l/2 \rfloor} \leq 1$.

6 Unary, binary and ternary Boolean predicates

It is enough to determine the values $\rho_l(P)$ for representatives of isomorphism classes of essentially unary, binary and ternary Boolean predicates. The case of 1-extendable Boolean predicates was handled in Theorem 3. The only essentially unary, binary and ternary Booleans predicates which are not 1-extendable (up to isomorphism) are the following: $P(x) = x$, $P(x, y) = x \wedge y$, $P(x, y, z) = x \wedge y \wedge z$, $P(x, y, z) = x \wedge (y \Leftrightarrow z)$ and $P(x, y, z) = x \wedge (y \vee z)$. The first three of these predicates satisfy that $\sigma(P) = 1$ and the values $\rho_l(P)$ are determined by Lemma 3. Therefore, we know the values $\rho_l(P)$ for all essentially unary and binary Boolean predicates. Tables 1 and 2 summarize our results. We focus on l -consistent sets of predicates for $P(x, y, z) = x \wedge (y \Leftrightarrow z)$ and $P(x, y, z) = x \wedge (y \vee z)$ in the rest of this section.

6.1 The predicate $P(x, y, z) = x \wedge (y \Leftrightarrow z)$

We first handle the case of 2-consistent sets:

Lemma 9 *It holds that $\rho_2(P) = \rho_2^w(P) = 8/27$ for $P(x, y, z) = x \wedge (y \Leftrightarrow z)$.*

Proof: We first show that $\rho_2(P) \geq \rho_2^w(P) \geq 8/27$. Let us consider a 2-consistent set Σ of weighted predicates. Since Σ is 2-consistent, Σ does not contain two predicates such that the first argument of one of them is x and the first argument of the other is $\neg x$. Therefore, we may assume that the first argument of each predicate is a positive literal.

Choose now each variable of Σ randomly and independently to be true with the probability $p = 2/3$. The probability that any single predicate of Σ is satisfied is either $p(p^2 + (1-p)^2) = 10/27$, if both the second and the third argument of it are positive or negative literals, or $4p^2(1-p) = 8/27$, otherwise. Hence, the expected fraction of satisfied constraints is at least $8/27$ and consequently $\rho(\Sigma) \geq 8/27$. Since the choice of a 2-consistent set Σ was arbitrary, we can conclude that $\rho_2(P) \geq 8/27$.

It remains to show that $\rho_2^w(P) \leq \rho_2(P) \leq 8/27$. For an integer $n \geq 3$, we consider a set Σ_n of predicates of type P with the variables x_1, \dots, x_n . Σ_n is formed by all the $n(n-1)(n-2)$ predicates $P(x_i, x_j, \neg x_k)$ for $1 \leq i, j, k \leq n$ where all i, j and k are mutually distinct. The set Σ is clearly 2-consistent. We now compute $\rho(\Sigma_n)$. Consider a truth assignment which assigns the true value to exactly n' variables of Σ_n . Then, the number of satisfied constraints is precisely $n'((n'-1)(n-n') + (n-n')(n'-1))$. Thus, we can conclude that (set $q = n'/n$):

$$\rho(\Sigma_n) \leq \max_{0 \leq q \leq 1} \frac{qn((qn-1)(n-qn) + (n-qn)(qn-1))}{n(n-1)(n-2)} =$$

$$\max_{0 \leq q \leq 1} 2q^2(1-q) + O\left(\frac{1}{n}\right) = \frac{8}{27} + O\left(\frac{1}{n}\right).$$

Hence, $\rho_2(P) \leq 8/27$ as claimed. □

We are now ready to fully analyze locally consistent sets of predicates for $P(x, y, z) = x \wedge (y \Leftrightarrow z)$:

Theorem 4 *If P is the predicate $P(x, y, z) = x \wedge (y \Leftrightarrow z)$, then the following holds for all $l \geq 1$:*

$$\rho_l(P) = \rho_l^w(P) = \begin{cases} 1/4 & \text{if } l = 1, \\ 8/27 & \text{if } l = 2, \\ 1/2 & \text{otherwise.} \end{cases}$$

Proof: It follows that $\rho_1(P) = \rho_1^w(P) = 1/4$ and $\rho_2(P) = \rho_2^w(P) = 8/27$ from Lemmas 2 and 9, respectively. Hence, we can focus only on the case $l \geq 3$ in the rest. First, we show that $\rho_l^w(P) \leq \rho_l(P) \leq 1/2$. Consider an l -consistent set Σ' of predicates $P'(y, z) = (y \Leftrightarrow z)$ and with $\rho(\Sigma') \leq 1/2 + \varepsilon$. Such a set Σ' exists because $\rho_l(P') = 1/2$ by Theorem 3. Let y_1, \dots, y_n be the variables contained in Σ' . We construct an l -consistent set Σ of predicates P with $\rho(\Sigma) = \rho(\Sigma')$. Introduce a new variable x and for each predicate $P'(y_i, y_j)$, $P'(\neg y_i, y_j)$, $P'(y_i, \neg y_j)$ and $P'(\neg y_i, \neg y_j)$ include to Σ a predicate $P(x, y_i, y_j)$, $P(x, \neg y_i, y_j)$, $P(x, y_i, \neg y_j)$ and $P(x, \neg y_i, \neg y_j)$, respectively. Since Σ' is l -consistent, the set Σ of predicates P is l -consistent, too. It is also not hard to see that $\rho(\Sigma) = \rho(\Sigma')$. Therefore, $\rho(\Sigma) \leq 1/2 + \varepsilon$ and $\rho_l(P) \leq 1/2$.

We now prove that $\rho_l(P) \geq \rho_l^w(P) \geq 1/2$ for $l \geq 3$. Let Σ be an l -consistent set of weighted predicates P . Let X be the set of variables which appear as the first argument in some of the predicates of Σ and Y the set consisting of the remaining variables. Since Σ is 2-consistent, we can assume that the first argument of each predicate is a positive literal. In addition, since Σ is 3-consistent it does not contain neither a predicate $P(x, x', \neg x'')$ nor a predicate $P(x, \neg x', x'')$ for some $x, x', x'' \in X$. Therefore, if we set each variable of X to be true, then each predicate of Σ is either satisfied (i.e., all its arguments are set and the predicate is true) or at least one of its arguments contains a variable from the set Y . Choose now each variable of Y randomly and independently to be true with the probability $1/2$. Each predicate, which was not satisfied by fixing the values of variables from the set X , is now satisfied with the probability $1/2$. Therefore, the expected weight of satisfied predicates is equal to the half of their total weight. Hence, $\rho(\Sigma) \geq 1/2$ and consequently $\rho_l(P) \geq \rho_l^w(P) \geq 1/2$. □

6.2 The predicate $P(x, y, z) = x \wedge (y \vee z)$

Before we analyze locally consistent sets of predicates for $P(x, y, z) = x \wedge (y \vee z)$, we consider 2-consistent and 3-consistent such sets:

Lemma 10 *It holds that $\rho_2(P) = \rho_2^w(P) = \sqrt{3}/9$ for $P(x, y, z) = x \wedge (y \vee z)$.*

Proof: We first show that $\rho_2(P) \geq \rho_2^w(P) \geq 2\sqrt{3}/9$. Let us consider a 2-consistent set Σ of weighted predicates. Since Σ is 2-consistent, Σ does not

contain two predicates such that the first argument of one of them is x and the first argument of the other is $\neg x$. Therefore, we may assume that the first argument of each predicate is a positive literal.

Choose now each variable of Σ randomly and independently to be true with the probability $p = 3^{-1/2} > 1/2$. The probability that any single predicate of Σ is satisfied is at least $p(1-p^2) = 2\sqrt{3}/9$. Hence, the expected fraction of constraints which are satisfied is at least $2\sqrt{3}/9$ and consequently $\rho(\Sigma) \geq 2\sqrt{3}/9$. Since the choice of a 2-consistent set Σ was arbitrary, we can conclude that $\rho_2(P) \geq \rho_2^w(P) \geq 2\sqrt{3}/9$.

It remains to show that $\rho_2^w(P) \leq \rho_2(P) \leq 2\sqrt{3}/9$. For an integer $n \geq 3$, we consider a set Σ_n of the predicates P with the variables x_1, \dots, x_n . Σ_n is formed by all the $n(n-1)(n-2)/2$ constraints $P(x_i, \neg x_j, \neg x_k)$ for $1 \leq i, j, k \leq n$, $i \neq j$, $i \neq k$ and $j < k$. The set Σ is clearly 2-consistent. We now compute $\rho(\Sigma_n)$. Consider a truth assignment which assigns the true value to exactly n' variables x_1, \dots, x_n . Then, the number of satisfied constraints is precisely the following $n'((n-n')(n'-1) + (n-n')(n-n'-1)/2)$. We can now conclude that (set $q = n'/n$):

$$\begin{aligned} \rho(\Sigma_n) &\leq \max_{0 \leq q \leq 1} \frac{qn((n-qn)(qn-1) + (n-qn)(n-qn-1)/2)}{n(n-1)(n-2)/2} = \\ &\max_{0 \leq q \leq 1} \frac{q(1-q)q + (1-q)^2/2}{1/2} + O\left(\frac{1}{n}\right) = \frac{2\sqrt{3}}{9} + O\left(\frac{1}{n}\right). \end{aligned}$$

Hence, $\rho_2(P) \leq 2\sqrt{3}/9$ as claimed. □

Lemma 11 *It holds that $\rho_3^w(P) \leq \rho_3(P) \leq 1/2$ for $P(x, y, z) = x \wedge (y \vee z)$.*

Proof: For each $\varepsilon > 0$, we construct a 3-consistent set Σ with $\rho(\Sigma) < 1/2 + \varepsilon$. Let n be an integer whose exact value will be chosen later. We construct a set Σ_n of predicates P with variables x_i for $1 \leq i \leq 2n+1$ and y_i^A for $1 \leq i \leq 2n+1$ where A ranges through all n -element subsets of $\{1, \dots, 2n+1\} \setminus \{i\}$. The set Σ_n consists of predicates $P(x_i, \neg x_j, y_i^A)$ for all $1 \leq i, j \leq 2n+1$, $i \neq j$ and $j \in A$ and predicates $P(x_i, \neg x_j, \neg y_i^A)$ for all $1 \leq i, j \leq 2n+1$, $i \neq j$ and $j \notin A$. In particular, the number of predicates contained in Σ_n is $(2n+1)2n \binom{2n}{n}$. Clearly, Σ_n is 3-consistent.

Let us consider a truth assignment which satisfies the most number of predicates. Let n' be the number of the variables x_1, \dots, x_{2n+1} with the

true value. By symmetry, we can assume that the values of the variables $x_1, \dots, x_{n'}$ are true and the values of $x_{n'+1}, \dots, x_n$ are false. Observe that all the predicates whose first argument is one of the literals $x_{n'+1}, \dots, x_n$ are false. In particular, if $n' \leq n$, then less than half of the predicates are satisfied. We focus on the case $n' > n$ in the rest of the proof.

Consider now an integer i , $1 \leq i \leq n'$, and an n -element subset A of the set $\{1, \dots, 2n+1\} \setminus \{i\}$. If $|A \cap (\{1, \dots, n'\} \setminus \{i\})| > (n'-1)/2$, then the truth assignment (because it is optimal) assigns y_i^A the true value and, otherwise, it assigns y_i^A the false value. Hence, the number of predicates, which contain y_i^A and which are satisfied, is $(2n - n' + 1) + \max\{|A \cap (\{1, \dots, n'\} \setminus \{i\})|, n' - |A \cap (\{1, \dots, n'\} \setminus \{i\})|\}$. For a fixed integer i , the number of n -element subsets A of $\{1, \dots, 2n+1\} \setminus \{i\}$ with $\max\{|A \cap (\{1, \dots, n'\} \setminus \{i\})|, n' - |A \cap (\{1, \dots, n'\} \setminus \{i\})|\} \geq (1 + \varepsilon)(n'-1)/2$ is at most the following:

$$\begin{aligned} & \sum_{k=0}^{\frac{(1-\varepsilon)(n'-1)}{2}} \binom{n'-1}{k} \binom{2n+1-n'}{n-k} + \sum_{k=\frac{(1+\varepsilon)(n'-1)}{2}}^{n'-1} \binom{n'-1}{k} \binom{2n+1-n'}{n-k} \leq \\ & \sum_{\substack{0 \leq k \leq \frac{(1-\varepsilon)(n'-1)}{2} \\ \frac{(1+\varepsilon)(n'-1)}{2} \leq k \leq n'-1}} \binom{n'-1}{k} 2^{2n+1-n'} \leq \\ & 2e^{-\frac{\varepsilon^2(n'-1)/2}{3}} 2^{n'-1} 2^{2n+1-n'} = 2^{2n+1} e^{-\frac{\varepsilon^2(n'-1)}{6}}. \end{aligned}$$

Hence, for a fixed i , the number of satisfied predicates whose first argument is x_i is at most the following (recall that $n+1 \leq n'$):

$$\begin{aligned} & \left(2n - n' + 1 + \frac{(1 + \varepsilon)(n' - 1)}{2}\right) \binom{2n}{n} + 2n 2^{2n+1} e^{-\frac{\varepsilon^2(n'-1)}{6}} \leq \\ & \left(2n - \frac{n' - 1}{2} + \frac{2n\varepsilon}{2}\right) \binom{2n}{n} + 2n 2^{2n+1} e^{-\frac{\varepsilon^2 n}{6}} \leq \\ & \left(2n - \frac{n' - 1}{2} + n\varepsilon\right) \binom{2n}{n} + 2n 2^{2n+1} e^{-\frac{\varepsilon^2 n}{6}}. \end{aligned}$$

Consequently, the fraction of satisfied predicates of Σ_n does not exceed:

$$\frac{n' \left(\left(2n - \frac{n'-1}{2} + n\varepsilon\right) \binom{2n}{n} + 2n 2^{2n+1} e^{-\frac{\varepsilon^2 n}{6}} \right)}{(2n+1) 2n \binom{2n}{n}} \leq$$

$$\frac{((2n+1)n + nn'\varepsilon) \binom{2n}{n} + 2nn'2^{2n+1}e^{-\frac{\varepsilon^2 n}{6}}}{(2n+1)2n \binom{2n}{n}} \leq$$

$$\frac{1}{2} + \frac{\varepsilon}{2} + \frac{2^{2n+1}e^{-\frac{\varepsilon^2 n}{6}}}{\binom{2n}{n}} \leq \frac{1}{2} + \frac{\varepsilon}{2} + 2(2n+1)e^{-\frac{\varepsilon^2 n}{6}}.$$

We now choose n to be an integer such that $2(2n+1)e^{-\frac{\varepsilon^2 n}{6}} \leq \varepsilon/2$. Then, each truth assignment with $n' > n$ satisfies at most the fraction of $1/2 + \varepsilon$ of the predicates of Σ_n . Hence, $\rho(\Sigma_n) \leq 1/2 + \varepsilon$ as desired. \square

We are now ready to determine the values $\rho_l(P)$ for $P(x, y, z) = x \wedge (y \vee z)$:

Theorem 5 *Let P be the predicate $P(x, y, z) = x \wedge (y \vee z)$. Then, the following holds for all $l \geq 1$:*

$$\rho_l(P) = \rho_l^w(P) = \begin{cases} 3/8 & \text{if } l = 1, \\ 2\sqrt{3}/9 & \text{if } l = 2, \\ \rho_{l-2}(\Pi_{2\text{-SAT}}) & \text{otherwise.} \end{cases}$$

Proof: The equalities $\rho_1(P) = \rho_1^w(P) = 3/8$ and $\rho_2(P) = \rho_2^w(P) = 2\sqrt{3}/9$ follow from Lemmas 2 and 10, respectively. We first prove that $\rho_l(P) \geq \rho_l^w(P) \geq \rho_{l-2}(\Pi_{2\text{-SAT}})$ for $l \geq 3$. Let Σ be an l -consistent set of weighted predicates P and let X be the set of variables of Σ which appear as the first argument in some predicates of Σ . Since Σ is 2-consistent, we can assume that all the first arguments of the predicates of Σ are positive literals. Let Y be the set of the remaining variables of Σ .

We construct an auxiliary $(l-2)$ -consistent 2-CNF formula Φ with the variables Y as follows. Since Σ is 3-consistent, it does not contain a predicate $P(x, \neg x', \neg x'')$ where $x', x'' \in X$. We now construct the formula Φ . For each predicate $P(x, \neg x', y)$ and each predicate $P(x, y, \neg x')$ of Σ with $x' \in X$, we include the clause (y) to Φ . Similarly, for each predicate $P(x, \neg x', \neg y)$ and each predicate $P(x, \neg y, \neg x')$ with $x' \in X$, we include the clause $(\neg y)$. For each predicate $P(x, y, y')$ with $y, y' \notin X$, we include the predicate $y \vee y'$ to Φ . We proceed analogously for predicates $P(x, \neg y, y')$, $P(x, y, \neg y')$ and $P(x, \neg y, \neg y')$. The weights of the clauses are equal to the weights of the corresponding predicates. If several same clauses are included to Φ , we replace them by a single clause whose weight is equal to the sum of the weights.

We claim that the formula Φ is $(l - 2)$ -consistent. If this is not the case, let Γ be the minimum inconsistent set of clauses of Φ . By Lemma 7, Γ contains at most two clauses of size one. We now find an inconsistent set Γ' of at most $|\Gamma| + 2$ predicates of Σ . For each clause of Γ of size two, include to Γ' the predicate of Σ corresponding to that clause. For each clause (y) , $(\neg y)$, of Γ , include to Γ' the predicate $P(x, y, \neg x')$, $P(x, \neg y, \neg x')$, respectively, which corresponds to that clause, together with one of the predicates of Σ whose first argument is x' . Since Γ contains at most two clauses of size two, $|\Gamma'| \leq |\Gamma| + 2 = l$. Moreover, since Γ is inconsistent, Γ' is also inconsistent. This contradicts the fact that Σ is l -consistent.

Since the formula Φ is $(l - 2)$ -consistent, there is a truth assignment which satisfies the fraction of $\rho(\Phi) \geq \rho_{l-2}(\Pi_{2\text{-SAT}})$ clauses of Φ . Extend this truth assignment to all the variables of Σ by assigning the true value to each variable $x \in X$. All the predicates of Σ whose arguments contain solely the variables from the set X are satisfied and, in addition, the fraction of $\rho(\Phi)$ of the remaining predicates are also satisfied. Therefore, $\rho(\Sigma) \geq \rho(\Phi) \geq \rho_{l-2}(\Pi_{2\text{-SAT}})$. Since the choice of a set Σ was arbitrary, we can conclude that $\rho_l(P) \geq \rho_l^w(P) \geq \rho_{l-2}(\Pi_{2\text{-SAT}})$.

It remains to prove that $\rho_l^w(P) \leq \rho_l(P) \leq \rho_{l-2}(\Pi_{2\text{-SAT}})$ for $l \geq 3$. If $l = 3$, the upper bound follows from Lemma 11. For $l \geq 4$ and $\varepsilon > 0$, fix an $(l - 2)$ -consistent 2-CNF formula Φ with $\rho(\Phi) \geq \rho_{l-2}(\Pi_{2\text{-SAT}}) + \varepsilon$ such that each minimal inconsistent set of at most l clauses contain two clauses of size one. Such a formula Φ exists by Lemma 8. Moreover, we can assume that each clause of size two is contained in Φ at most once. Let m' be the number of clauses of Φ of size one (counting multiplicities) and m the number of all clauses of Φ . Since Φ is 2-consistent, $m'/m \leq \rho(\Phi)$. We now construct an l -consistent set Σ of predicates P with $\rho(\Sigma) = \rho(\Phi)$.

Let y_1, \dots, y_n be the variables contained in the formula Φ . The set Σ will contain $(m + 1)n$ variables y_i^j for $1 \leq i \leq n$ and $1 \leq j \leq m + 1$ and $m + 1$ variables x^j for $1 \leq j \leq m + 1$. Let C_1, \dots, C_m be the clauses of Φ . For each clause $C_k = (y_i \vee y_{i'})$, $1 \leq k \leq m$, we include to Σ predicates $P(x^j, y_i, y_{i'})$ for $1 \leq j \leq m + 1$. Similarly, we proceed for clauses $C_k = (y_i \vee \neg y_{i'})$ and $C_k = (\neg y_i \vee \neg y_{i'})$. If the clause C_k is of size one, say $C_k = (y_i)$, we include to Σ predicates $P(x^j, y_i, \neg x^{(j+k) \bmod (m+1)})$ for $1 \leq j \leq m + 1$. Therefore, Σ consists of $k(m + 1)$ distinct predicates.

First, we show that Σ is l -consistent. Assume the opposite and let Γ be the minimum inconsistent set of predicates contained in Σ , i.e., $|\Gamma| \leq l$. Observe that if we set all the variables x^1, \dots, x^n to be true, then Σ reduces

to $m + 1$ independent “copies” of the formula Φ . Therefore, if Γ is a set of l inconsistent predicates, it must contain predicates contained in one of these copies of Φ which correspond to an inconsistent set Γ_Φ of clauses of Φ . By symmetry, we can assume that predicates corresponding to Γ_Φ are contained in the first copy of Φ . Since Φ is $(l - 2)$ -consistent, $|\Gamma_\Phi| \geq l - 1$. On the other hand, $|\Gamma_\Phi| \leq |\Gamma| \leq l$. By the choice of Φ , each inconsistent set of at most l clauses of Φ contains two clauses of size one. Let $C_k = (y_i)$ and $C_{k'} = (y_{i'})$ be these two clauses of size one, i.e., Γ contains the predicates $P(x^1, y_i, \neg x^{(k+1) \bmod (m+1)})$ and $P(x^1, y_{i'}, \neg x^{(k'+1) \bmod (m+1)})$. If Γ is inconsistent, it must contain a predicate whose first argument is $x^{(k+1) \bmod (m+1)}$ as well as a predicate whose first argument is $x^{(k'+1) \bmod (m+1)}$. Therefore, Γ contains at least $|\Gamma_\Phi| + 2 > l$ predicates. Hence, the set Σ is l -consistent.

It remains to show that $\rho(\Sigma) = \rho(\Phi)$. Since $\rho(\Phi) \leq \rho_{l-2}(\Pi_{2\text{-SAT}}) + \varepsilon$ and the choice of ε was arbitrary, this would yield $\rho_l(P) \leq \rho_{l-2}(\Pi_{2\text{-SAT}})$. Fix a truth assignment such that the fraction of $\rho(\Sigma)$ predicates of the set Σ is satisfied. We claim that there is a truth assignment which assigns all the variables x^1, \dots, x^{m+1} the true value. Indeed, if x^j is false, then change of its value to true. This causes at most m' previously satisfied predicates to be unsatisfied (precisely those which contain $\neg x^j$ as the third argument) and, on the other hand, we can choose values of y_1^j, \dots, y_n^j so that at least the $\rho(\Phi)m$ predicates whose first argument is x^j are satisfied. Note that none of these $\rho(\Phi)m$ predicates could be satisfied before the change. Since $\rho(\Phi)m \geq m'$ (recall that $\rho(\Phi) \geq m'/m$), the number of satisfied predicates is not decreased after the change. Hence, we assume that all the variables x^1, \dots, x^{m+1} are set to be true by the considered truth assignment. Then, the set Σ is reduced to $m + 1$ independent “copies” of the formula Φ (substitute the true value for all the variables x^1, \dots, x^{m+1}). We can conclude that $\rho(\Sigma) = \rho(\Phi)$. \square

7 Conclusion

We studied instances of constraint satisfaction problems which are locally consistent. The values of $\rho_l(P)$ have been determined for all predicates which have arity at most three or which are 1-extendable. We were not able to fully analyze non-1-extendable predicates. Already, the smallest two non-trivial such predicates, namely $P(x, y, z) = x \wedge (y \Leftrightarrow z)$ and $P(x, y, z) = x \wedge (y \vee z)$,

showed that the behavior of locally consistent P -systems for such predicates P can be quite weird.

For the asymptotic behavior of $\rho_\infty(\Pi)$, we settled almost completely the case of finite sets Π of predicates. The only case which remains open is to determine $\rho_\infty(\Pi)$ for sets of predicates Π which contain a predicate of arity one. The case of infinite sets Π seems to be also interesting, but rather from the theoretical point of view than the algorithmic one: in most cases, it might be difficult to describe the input if the set Π is not a “nice” set of predicates as it is the case of, e.g., Π_{SAT} . For an infinite set of predicates Π , one can also define the set $\pi(\Pi)$ and then $\Psi(\pi(\Pi))$ to be the infimum of Ψ taken over all convex combinations of finite number of functions from $\pi(\Pi)$. It is not hard to verify that the proof of Theorem 2 can be translated to this setting. In particular, $\rho_\infty^w(\Pi) \geq \Psi(\pi(\Pi))$ for every infinite set Π . However, the proof of Theorem 1 cannot be adopted to this case since the arity of the predicates of Π is not bounded. We suspect that the equality $\rho_\infty^w(\Pi) = \Psi(\pi(\Pi))$ does not hold for all (infinite) sets Π .

Acknowledgement

The authors would like to thank Gerhard Woeginger for attracting their attention to locally consistent formulas and for pointing out several useful references for the problem. They would also like to thank Dimitrios M. Thilikos for suggesting the version of the problem considered in this paper.

References

- [1] N. Alon, J. Spencer: *The probabilistic method*, 2nd edition, John Wiley, New York, 2000.
- [2] B. Chazelle: *The discrepancy method: Randomness and complexity*, Cambridge University Press, 2000.
- [3] S. Cook, The complexity of theorem-proving procedures, Proc. of the 3rd ACM Symposium on Theory of Computing, ACM, New York (1971) 29–33.
- [4] S. Cook, D. Mitchell, Finding Hard Instances of the Satisfiability Problem: A Survey, In: Satisfiability Problem: Theory and Applications,

DIMACS Series in Discrete Mathematics and Theoretical Computer Science **35** AMS (1997).

- [5] D. Eppstein, Improved Algorithms for 3-coloring, 3-edge-coloring and Constraint Satisfaction, In: Proc. of the 12th ACM-SIAM Symposium on Discrete Algorithms, SIAM (2001) 329–337.
- [6] T. Feder, R. Motwani, Worst-case Time Bounds for Coloring and Satisfiability Problems, J. Algorithms **45(2)** (2002) 192–201.
- [7] T. Hagerup, Ch. Rüb, A guided tour Chernoff bounds, Inform. Process. Letters **33** (1989) 305–308.
- [8] J. Håstad, Some optimal inapproximability results, Journal of ACM **48(4)** (2001) 798–859. A preliminary version appeared in: Proc. 28th ACM Symposium on Theory of Computing (STOC), ACM Press (1997) 1–10.
- [9] M. A. Huang, K. Lieberherr, Implications of Forbidden Structures for Extremal Algorithmic Problems, Theoretical Computer Science **40** (1985) 195–210.
- [10] S. Jukna: *Extremal Combinatorics with Applications in Computer Science*, Springer, Heidelberg, 2001.
- [11] D. Král', Locally Satisfiable Formulas, In: Proc. of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM (2004) 323–332.
- [12] K. Lieberherr, E. Specker, Complexity of Partial Satisfaction, J. of the ACM **28(2)** (1981) 411–422.
- [13] K. Lieberherr, E. Specker, Complexity of Partial Satisfaction II, Technical Report 293, Dept. of EECS, Princeton University (1982).
- [14] R. Motwani, P. Raghavan: *Randomized Algorithms*, Cambridge University Press, 1995.
- [15] T. J. Schaefer, The complexity of satisfiability problems, In: Proc. of the 10th Annual ACM Symposium on Theory of Computing (STOC), ACM Press (1978) 216–226.

- [16] L. Trevisan, On Local versus Global Satisfiability, *SIAM J. Disc. Math.* (to appear). A preliminary version is available as ECCC report TR97-12.
- [17] Z. Usiskin, Max-min Probabilities in the Voting Paradox, *Ann. Math. Stat.* **35** (1963) 857–862.
- [18] G. J. Woeginger, Exact Algorithms for NP-hard Problems: A Survey, In: M. Jünger, G. Reinelt, G. Rinaldi (eds.), *Proc. 5th Int. Worksh. Combinatorial Optimization - Eureka, You Shrink*, Lecture Notes in Computer Science **2570**, Springer-Verlag Berlin (2003) 185–207.
- [19] M. Yannakakis, On the Approximation of Maximum Satisfiability, *J. Algorithms* **17** (1994) 475–502. A preliminary version appeared in: *Proc. of the 3rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM (1992) 1–9.