An Asymptotically Optimal Linear-Time Algorithm for Locally Consistent Constraint Satisfaction Problems

Daniel Kráľ   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. We study the asymptotic behavior of $\rho_l(\Pi)$ for sets Π consisting of Boolean predicates. The value $\rho_\infty(\Pi) := \lim_{l\to\infty} \rho_l(\Pi)$ is determined for all such sets Π. Moreover, we design a robust deterministic algorithm (for a fixed set Π of predicates) 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. Most of our results hold for both the unweighted and weighted versions of the problem.

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

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

satisfaction problems [3, 5, 6, 18]. However, some instances of real problems do not require all the constraints to be satisfied but it is enough to satisfy a large fraction of them. In order to maximize this fraction, the input can be usually pruned at the beginning by removing small sets of contradictory constraints so the input instance is usually "locally" consistent. Formally, an instance of the (weighted) constraint satisfaction problem is *l-consistent* if any l constraints of it can be simultaneously satisfied. In a weighted version of the problem the constraints are assigned (positive) weights and the goal is to maximize the total weight of satisfied constraints. In this paper, we design a robust linear-time algorithm for l-consistent constraint satisfaction problems whose constraints are Boolean predicates which is asymptotically optimal as l tends to infinity.

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. Similarly, $\rho_l^w(\Pi)$ denotes this maximum for the weighted version of the problem (see Section 2 for more formal definitions). Let further $\rho_\infty(\Pi) = \lim_{l\to\infty} \rho_l(\Pi)$ and $\rho_\infty^w(\Pi) = \lim_{l\to\infty} \rho_l^w(\Pi)$. 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 8). We postpone the formal definitions of the functional Ψ and the set $\pi(\Pi)$ to Section 2. Examples how to apply this result can be found in Examples 3 and 4. Some of our results also hold for the case when the set Π is infinite as discussed in Section 5.

The main algorithmic result of this paper (Theorem 2) 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, 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.

1.1 Previous results and their relation to our results

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 GF(2). However, even if this 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 [4]. In particular, $\rho_\infty(\Pi) = 1/2$ and our algorithm achieves the best possible ratio.

Locally consistent constraint satisfaction problems for constraints which are Boolean predicates was first studied by Trevisan [16] and since then they have attracted a substantial interest of researchers. Trevisan [16] proved that if Π is the set of all the Boolean predicates of arity k, then $\rho^w_\infty(\Pi) = \rho_\infty(\Pi) = 2^{1-k}$. Dvořák et al. [4] showed that if Π is a set containing a single 1-extendable Boolean predicate P of arity k (see Section 2 for the definition of 1-extendibility), then $\rho^w_l(\Pi) = \rho_l(\Pi) = \sigma(P)/2^k$ for all $l \geq 1$ where $\sigma(P)$ is the number of possible combinations of arguments which satisfy the predicate P. In particular, $\rho^w_\infty(\Pi) = \rho_\infty(\Pi) = \sigma(P)/2^k$. In [4], all the values $\rho^w_l(\Pi)$ has also been determined for sets Π consisting of a single Boolean predicate with the arity $k \leq 3$, e.g., it was shown in [4] that $\rho^w_\infty(\Pi^k) = 3/4$ for $k = 3$ where Π^k is the set containing a single (non-1-extendable) predicate $P(x_1, \ldots, x_k) = x_1 \wedge (x_2 \vee \cdots \vee x_k)$. Our results imply that $\rho^w_\infty(\Pi^k) = 3/4$ for all $3 \leq k \leq 6$ (see Examples 2 and 4). However, surprisingly, $\rho^w_\infty(\Pi^k) > 3/4$ for all $k \geq 7$ as shown in Example 5.

The most studied variant of the 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, $\Pi_{2-\mathrm{SAT}}$ 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 $\rho_l^w(\Pi_{\mathrm{SAT}})$ and $\rho_l^w(\Pi_{2-\mathrm{SAT}})$ are known only for small values of l: clearly, $\rho_1^w(\Pi_{\mathrm{SAT}}) = \rho_1^w(\Pi_{2-\mathrm{SAT}}) = 1/2$. Lieberherr and Specker [12] showed that $\rho_2^w(\Pi_{\mathrm{SAT}}) = \rho_2^w(\Pi_{2-\mathrm{SAT}}) = \frac{\sqrt{5}-1}{2} \approx 0.6180$ and subsequently [13] they showed that $\rho_3^w(\Pi_{\mathrm{SAT}}) = \rho_3^w(\Pi_{2-\mathrm{SAT}}) = 2/3$. 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, $\rho_4^w(\Pi_{\mathrm{SAT}}) \approx 0.6992$ but $\rho_4^w(\Pi_{2-\mathrm{SAT}}) > 0.6992$. Second, the values $\rho_l^w(\Pi_{\mathrm{SAT}})$ 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 $\rho_4^w(\Pi_{\mathrm{SAT}})$ differs from the value 0.6920 for the fractional version of the problem.

The asymptotic behavior of $\rho_l^w(\Pi_{\mathrm{SAT}})$ was first addressed by Huang and Lieberherr [9] who showed that $\rho_\infty^w(\Pi_{\mathrm{SAT}}) \leq 3/4$. The limit was settled by Trevisan [16] who showed $\rho_\infty^w(\Pi_{\mathrm{SAT}}) = 3/4$. Trevisan's result also yields that $\rho_\infty^w(\Pi_{2-\mathrm{SAT}}) = 3/4$. 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

In the 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. This does not decrease generality of our results: if a single variable is allowed to be contained in several distinct arguments of a single predicate, 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).

A Boolean predicate P is 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. 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 $\pi_{P,\tau}(p) : \langle 0, 1\rangle \to \langle 0, 1\rangle$ be equal to the probability that the k-ary predicate P with arguments $x_1, \ldots, x_k$ 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 $P_2(x_1, x_2) = (x_1 \wedge x_2)$. 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:*

$$\pi_{P_1,1}(p) = p \qquad \pi_{P_2,11}(p) = 2p - p^2$$

$$\pi_{P_2,1\star}(p) = \pi_{P_2,\star 1}(p) = (p+1)/2 \qquad \pi_{P_2,\star\star}(p) = 3/4.$$

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 $f : \langle 0, 1\rangle \to \langle 0, 1\rangle$ its maximum on the interval $\langle 0, 1\rangle$. If F is a finite family of functions $f : \langle 0, 1\rangle \to \langle 0, 1\rangle$, 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 $\rho_\infty(\Pi) = \lim_{l\to\infty} \rho_l(\Pi)$ is equal to $\Psi(\pi(\Pi))$ for any set Π of Boolean predicates with arities at least two and $\rho^w_\infty(\Pi)$ is equal to $\Psi(\pi(\Pi))$ for any set Π of Boolean predicates (see Corollary 8 and Examples 3–5 after it).

3 The algorithm and the upper bound

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

Lemma 1 *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:

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

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

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

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

$$|f'(p)| \le \sum_{j_0=1}^{k'} \left| \sum_{i_{j_0}=0}^{1} (-1)^{1+i_{j_0}} \sum_{i_1=0}^{1} \cdots \sum_{i_{j_0-1}=0}^{1} \sum_{i_{j_0+1}=0}^{1} \cdots \sum_{i_{k'}=0}^{1} \alpha_{i_1,\ldots,i_{k'}} \prod_{j=1,j\neq j_0}^{k'} f_{i_j}(p) \right|$$

$$\leq \sum_{j_0=1}^{k'} 1 \;=\; k' \;\leq\; K$$

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 $j_0 = 1, \ldots, k'$. Since both the function f_0 and f_1 are non-negative, the value of the function

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

is always between 0 and 1 for $p \in \langle 0, 1\rangle$, $j_0 = 1, \ldots, k'$ 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 2 *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 of 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 $\mu_0, \ldots, \mu_{\lceil 2K/\varepsilon \rceil}$ and subsets $\Sigma_1, \ldots, \Sigma_{\lceil 2K/\varepsilon \rceil}$ 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 $\mu_0, \ldots, \mu_{i-1}$ 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 $\Sigma_1, \ldots, \Sigma_{\lceil 2K/\varepsilon \rceil}$.

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:

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

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:

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

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 $|A|K \le |\Sigma|K$.

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 $\Sigma_1, \ldots, \Sigma_{\lceil 2K/\varepsilon \rceil}$, the total weight of all the predicates contained in the sets $\Sigma_1, \ldots, \Sigma_{\lceil 2K/\varepsilon \rceil}$ when counting multiplicities does not exceed Kw_0. By an averaging argument, there exists $1 \le i \le \lceil 2K/\varepsilon \rceil$ 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 $w_0' \ge (1 - \varepsilon/2)w_0$ 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 1.

Compute the value of the function $f(p)$ for each of the following values of p: $0, \frac{\varepsilon}{K}, \frac{2\varepsilon}{K}, \ldots, \left\lfloor \frac{K}{\varepsilon} \right\rfloor \frac{\varepsilon}{K}, 1$. 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 2 is the following:

Corollary 3 *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 The 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 4 *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 $i = 1, \ldots, l$ 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 $P_1, \ldots, P_i$ be any i predicates of Σ. Since $G(\Sigma')$ contains no cycle of length at most l, the vertices corresponding to $P_1, \ldots, P_i$ 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 $y_1, \ldots, y_n$ 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 $y_1, \ldots, y_n$ which satisfies all the predicates $P_1, \ldots, P_{i-1}$. Since P_i is a leaf or an isolated vertex in T, it has at most one variable in common with the predicates $P_1, \ldots, P_{i-1}$. Hence, the truth assignment for $y_1, \ldots, y_n$ can be extended to a truth assignment which satisfies all the predicates $P_1, \ldots, P_i$ 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 5 *Let X be a non-negative random variable with the expected value E. The following holds for every $\alpha \geq 1$:*

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

Proposition 6 *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 $0 < \delta \leq 1$:*

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

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

Theorem 7 *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 $f_1, \ldots, f_K$ 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 $x_1, \ldots, x_n$ and $y_1, \ldots, y_n$ where n is a sufficiently large power of two which will be fixed later in the proof. Fix an integer $i = 1, \ldots, K$ 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 $n^k 2^{k'}$ predicates P^i whose j-th argument, $1 \leq j \leq k$, is a positive literal containing one of the variables $x_1, \ldots, x_n$ if $\tau^i_j = 1$, a negative literal containing one of the variables $x_1, \ldots, x_n$ if $\tau^i_j = 0$ and a positive or negative literal containing one of the variables $y_1, \ldots, y_n$ if $\tau^i_j = \star$, 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 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 system 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 6 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 $\max\{f_i(n'/n), \varepsilon/8\}n^k 2^{k'} \cdot \alpha_i n^{-(k-1)+1/2l} 2^{-k'}$. Using the fact that $f_i(n'/n) \le 1$ and Proposition 6, we infer the following:

$$\text{Prob}(\mu \text{ satisfies more than } \alpha_i(f_i(n'/n) + \frac{\varepsilon}{4})n^{1+1/2l} \text{ predicates of } \Sigma^i) \le$$

$$\text{Prob}(\Sigma^i \text{ contains more than } \alpha_i(f_i(n'/n) + \tfrac{\varepsilon}{4})n^{1+1/2l} \text{ good predicates}) \le$$

$$\text{Prob}(\Sigma^i \text{ contains } > (1+\varepsilon/8)\alpha_i \max\{f_i(n'/n), \varepsilon/8\}n^{1+1/2l} \text{ good predicates}) \le$$

$$e^{-\frac{\varepsilon^2 \alpha_i \max\{f_i(n'/n),\varepsilon/8\}n^{1+1/2l}}{192}} \le e^{-\frac{\varepsilon^3 \alpha_i n^{1+1/2l}}{1536}}$$

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 $n^k 2^{k'}$ 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 $\binom{k}{2} n^{k-1} 2^{k'} \alpha_i 2^{-k'} n^{-(k-1)+1/2l} = \alpha_i \binom{k}{2} n^{1/2l}$. By Markov's inequality (Proposition 5), 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 $x_1, \ldots, x_n$ 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 $x_1, \ldots, x_n$ 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 $P'_1, \ldots, P'_{l'}$, $2 \le l' \le l$, described by $\tau'_1, \ldots, \tau'_{l'}$. 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 $P'_1, \ldots, P'_{l'}$ determined by $\tau'_1, \ldots, \tau'_{l'}$ 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 $P'_1, \ldots, P'_{l'}$ determined by $\tau'_1, \ldots, \tau'_{l'}$ is at most $\prod_{i=1}^{l'} k'_i n^{1/2l} \le K_0^{l'} n^{1/2}$ (recall that $0 \le \alpha'_i \le 1$ for all $1 \le i \le l'$ 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 $P'_1, \ldots, P'_{l'}$ and $3^{K_0 l'}$ possible choices of the vectors $\tau'_1, \ldots, \tau'_{l'}$, the expected number of semicycles of Σ of length l' does not exceed $(KK_0 3^{K_0})^{l'} n^{1/2}$. By Proposition 5, 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:

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

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 $y_1, \ldots, y_n$, 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' $(0 \le n' \le n)$ of the variables $x_1, \ldots, x_n$ to be true:

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

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

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 4. Consequently, $\rho_l^w(\Pi) \le \Psi(\pi(\Pi)) + \varepsilon$. Moreover, if the arity of each predicate of Π is at least two, the weights of all the predicates of Σ are one and $\rho_l(\Pi) \le \Psi(\pi(\Pi)) + \varepsilon$.

■

We immediately infer from Corollary 3 and Theorem 7 the following expressions for $\rho_\infty(\Pi)$ and $\rho_\infty^w(\Pi)$:

Corollary 8 *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 8, 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 $0 \le p \le 1$, we infer $\Psi(\pi(\Pi)) \le \Psi(\pi_{P_2,\star\star}) = 3/4$. 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 $\Psi(\pi(\Pi)) \ge 3/4$. 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* $\Psi(\pi(\Pi)) \le \Psi(\pi_{P,100\star\star}) = 3/4$*. 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,* $\rho^w_\infty(\Pi) = \Psi(\pi(\Pi)) = 3/4$*.*

Example 5 *Let* Π^k *be the set containing a single predicate* $P(x_1, \ldots, x_k) = x_1 \wedge (x_2 \vee \cdots \vee x_k)$ *for an integer* $k \ge 7$*. Consider the vector* $\tau = 10\cdots0\star\star$*. 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* $\rho^w_\infty(\Pi^k) = \Psi(\pi(\Pi^k)) \ge \Psi(\pi_{P,\tau}) > 3/4$*.*

5 Conclusion

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 contains 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 7 can be translated to this setting. In particular, $\rho^w_\infty(\Pi) \ge \Psi(\pi(\Pi))$ for every infinite set Π. However, the proof of Theorem 2 cannot be adopted to this case since the arity of the predicates of Π is not bounded. We suspect that the equality $\rho^w_\infty(\Pi) = \Psi(\pi(\Pi))$ 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, 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, Vol. 35 AMS (1997).

[4] Z. Dvořák, D. Kráľ, O. Pangrác: Locally consistent constraint satisfaction problems. To appear in: Proc. of 31st International Colloquium on Automata, Languages and Programming (ICALP), 2004.

[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 (2001). 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. Kráľ: 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, Vol. 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.