

# Parameterized Complexity of DPLL Search Procedures

OLAF BEYERSDORFF, University of Leeds and Sapienza University of Rome  
NICOLA GALESÌ, Sapienza University of Rome  
MASSIMO LAURIA, KTH Royal Institute of Technology

We study the performance of DPLL algorithms on parameterized problems. In particular, we investigate how difficult it is to decide whether small solutions exist for satisfiability and other combinatorial problems. For this purpose we develop a Prover-Delayer game that models the running time of DPLL procedures and we establish an information-theoretic method to obtain lower bounds to the running time of parameterized DPLL procedures. We illustrate this technique by showing lower bounds to the parameterized pigeonhole principle and to the ordering principle. As our main application we study the DPLL procedure for the problem of deciding whether a graph has a small clique. We show that proving the absence of a  $k$ -clique requires  $n^{\Omega(k)}$  steps for a nontrivial distribution of graphs close to the critical threshold. For the restricted case of tree-like Parameterized Resolution, this result answers a question asked by Beyersdorff et al. [2012] of understanding the Resolution complexity of this family of formulas.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—Complexity of proof procedures

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Proof complexity, parameterized complexity, resolution, prover-delayer games

## ACM Reference Format:

Beyersdorff, O., Galesi, N., and Lauria M. 2013. Parameterized complexity of DPLL search procedures. *ACM Trans. Comput. Logic* 14, 3, Article 20 (August 2013), 21 pages.  
DOI: <http://dx.doi.org/10.1145/2499937.2499941>

## 1. INTRODUCTION

Resolution was introduced by Blake [1937] and since the work of Robinson [1965] and Davis, Putnam, Logemann, and Loveland [Davis and Putnam 1960; Davis et al. 1962] has been highly employed in proof search and automated theorem proving. In the last years, the study of Resolution has gained great significance in at least two important fields of computer science. (1) *Proof complexity*, where Resolution is one of the most intensively investigated proof systems [Alekhovich and Razborov 2008; Beame and Pitassi 1996; Ben-Sasson and Wigderson 2001; Bonnet et al. 2000; Chvátal and Szemerédi 1988; Haken 1985; Urquhart 1987]. The study of lower bounds for proof

---

A preliminary version of the results in this article appeared in the proceedings of SAT'11 [Beyersdorff et al. 2011].

This research was supported by the grant “Limits of Theorem Proving” from the John Templeton Foundation and a DAAD grant.

Authors' addresses: O. Beyersdorff, School of Computing, University of Leeds; email: [o.beyersdorff@leeds.ac.uk](mailto:o.beyersdorff@leeds.ac.uk); N. Galesi, Dipartimento di Informatica, Sapienza Università di Roma; M. Lauria, School of Computer Science and Communication, KTH Royal Institute of Technology.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2013 ACM 1529-3785/2013/08-ART20 \$15.00

DOI: <http://dx.doi.org/10.1145/2499937.2499941>

length in this system has opened the way to lower bounds in much stronger proof systems [Beame et al. 1996; Segerlind et al. 2004]. (2) *Algorithms for the satisfiability problem* of CNF formulas, where the DPLL algorithm [Beame et al. 2002; Davis et al. 1962] was at the core of the most important algorithms employed for the satisfiability problem [Beame et al. 2002], until the discovery of *conflict driven clause learning* solvers [Beame et al. 2004]. Running DPLL on unsatisfiable formulas produces Resolution refutations in the simple form of a tree, thus Resolution proof lengths are connected with the running time of DPLL procedures.

*Parameterized Resolution* was recently introduced by Dantchev et al. [2011] in the context of *parameterized proof complexity*, an extension of the proof complexity approach of Cook and Reckhow [1979] to parameterized complexity. Analogously to the case of Fixed Parameter Tractable (FPT) algorithms for optimization problems, the study of Parameterized Resolution provides new approaches and insights to proof search and to proof complexity. Loosely speaking, to refute a parameterized contradiction  $(F, k)$  in Parameterized Resolution we have built-in access to new *axioms*, which encode some property on assignments. In the most common case the new axioms are the clauses forbidding assignments of hamming weight greater than  $k$ . We underline that only those axioms appearing in the proof account for the proof length. Hence Parameterized DPLL refutations can be viewed as traces of executions of a (standard) DPLL algorithm in which some branches are cut because they falsify one of the new axioms.

In spite of its recent introduction, research in this direction is already active. Gao [2009] analyzes the effect of the standard DPLL algorithm on the problem of weighted satisfiability for random  $d$ -CNFs. Beyersdorff et al. [2012], using an idea also developed by Chen and Flum [2008], proved that there are FPT efficient Parameterized Resolution proofs for *all* bounded-width unsatisfiable CNF formulae.

As our first contribution, we look inside the structure of Parameterized DPLL giving a new information-theoretical characterization of proofs in terms of a two-player game, the *Asymmetric Prover-Delayer (APD) game*. The APD-game was also used by Beyersdorff et al. [2010b] to prove simplified optimal lower bounds for the pigeonhole principle in tree-like classical Resolution. Compared to Beyersdorff et al. [2010b], we present here a completely different analysis of APD-games based on an information-theoretical argument that is new and interesting by itself.

Parameterized Resolution is also a refutational proof system for parameterized contradictions. Hence proving proof length lower bounds for parameterized contradictions is important in order to understand the strength of such a proof system. Dantchev et al. [2011] proved significant lower bounds for Parameterized DPLL proofs of the pigeonhole principle (PHP) and of the ordering principle (OP). Moreover, recently the work [Beyersdorff et al. 2012] extended the PHP lower bounds to the case of parameterized dag-like bounded-depth Frege.<sup>1</sup>

As our second contribution we provide a unified approach to get lower bounds in Parameterized DPLL using the APD-game. As a simple application of our characterization, we give an improvement on the lower bounds given by Dantchev et al. [2011] for both PHP and OP.

It is a natural question what happens when we equip a proof system with a more efficient way of encoding the exclusion of assignments with hamming weight  $\geq k$ , rather than just adding all possible clauses with  $k + 1$  negated variables. Dantchev

<sup>1</sup>The APD-game appeared also in the technical report Beyersdorff et al. [2010a], together with a lower bound for dag-like Parameterized Resolution, but all results in Beyersdorff et al. [2010a] are subsumed and improved by Beyersdorff et al. [2012] and the present article.

et al. [2011] proved that this is a relevant issue. They presented a different and more efficient encoding, and showed that under this encoding PHP admits efficient FPT Parameterized Resolution proofs.

Beyersdorff et al. [2012] investigate this question further and notice that for propositional encodings of prominent combinatorial problems like  $k$ -independent set or  $k$ -clique, the separation between the two encodings vanishes. Hence they proposed (see Question 5 in Beyersdorff et al. [2012]) to study the performance of Parameterized Resolution on CNF encodings of such combinatorial problems and in particular to prove lower bounds. This will capture the real proof-theoretic strength of Parameterized Resolution, since it is independent of the encodings. The  $k$ -clique principle (see also Beyersdorff et al. [2012] and Beame et al. [2007] for similar principles) simply says that a given graph contains a clique of size  $k$ . When applied on a graph not containing a  $k$ -clique it is a contradiction.

As a third contribution, we prove significant lower bounds for the  $k$ -clique principle in the case of Parameterized DPLL. Our  $k$ -clique formula is based on random graphs distributed according to a simple variation of the Erdős-Rényi model  $G(n, p)$ . It is well known [Janson et al. 2000, Chapter 3] that when  $G$  is drawn according to  $G(n, p)$  and  $p \ll n^{-\frac{2}{k-1}}$ , with high probability  $G$  has no  $k$ -clique. We introduce a canonical CNF *Clique*( $G, k$ ) expressing this fact and show that with high probability these formulas are hard for Parameterized DPLL.

For the canonical graphs with no  $k$ -clique, that is, the  $(k-1)$ -partite complete graph, we show that the same principle admits fixed parameterized tractable refutations in dag-like Resolution, but not in tree-like. As an open problem it remains whether this is the case also for the random graphs above.

The article is organized as follows. Section 2 contains all preliminary notions and definitions concerning fixed-parameter tractability, parameterized proof systems, and Parameterized Resolution. In Section 3 we define the asymmetric Prover-Delayer game and establish its precise relation to the proof size in tree-like Parameterized Resolution. In Section 4 we show example applications for the APD-game: we give tree-like Parameterized Resolution lower bounds for the pigeonhole principle and for an ordering principle. Our bounds improve the ones that follows from the model-theoretic criteria in Dantchev et al. [2011].

In Section 5 we introduce the formula *Clique*( $G, k$ ), which is satisfiable if and only if there is a  $k$ -clique in the graph  $G$  and we show that on a certain distribution of random graphs the following holds with high probability:  $G$  has no  $k$ -clique and the size of the shortest refutation of *Clique*( $G, k$ ) is  $n^{\Omega(k)}$ . From an algorithmic perspective, this result can be formulated as: any algorithm for  $k$ -clique that (i) cleverly selects a vertex and branches on whether it is in the clique or not, (ii) deletes all its nonneighbors, and (iii) stops branching when there are no vertices left, must use at least  $n^{\Omega(k)}$  steps for most random graphs with a certain edge probability.

## 2. PRELIMINARIES

*Parameterized complexity* is a branch of complexity theory where problems are analyzed in a finer way than in the classical approach: instead of expressing the complexity of a problem as a function only of the input size there is an extra parameter that is part of the input, and one investigates the effect of the parameter on the complexity. We say that a problem is *fixed-parameter tractable* (FPT) with parameter  $k$  if it can be solved in time  $f(k)n^{O(1)}$  for some computable function  $f$  of arbitrary growth. In this setting classically intractable problems may have efficient solutions, assuming the parameter is small, even if the total size of the input is large. Parameterized

complexity also has a completeness theory: many parameterized problems that appear not to be fixed-parameter tractable have been classified as being complete under fixed-parameter tractable reductions for the complexity classes in the so-called weft hierarchy  $W[1] \subseteq W[2] \subseteq W[3] \subseteq \dots$ . Parameterized complexity is a rich and informative theory, and we suggest the monographs [Downey and Fellows 1999; Flum and Grohe 2006; Niedermeier 2006] for further reading about FPT and the weft hierarchy.

Consider the problem BOUNDED CNF SAT of deciding whether there is a satisfying assignment of Hamming weight at most  $k$  for a formula in conjunctive normal form.<sup>2</sup> Many parameterized combinatorial problems can be naturally encoded in BOUNDED CNF SAT: finding a vertex cover of size at most  $k$ ; finding a clique of size at least  $k$ ; or finding a dominating set of size at most  $k$ . In the theory of parameterized complexity, the hardness of the BOUNDED CNF SAT problem is reflected by the fact that it is  $W[2]$ -complete [Beyersdorff et al. 2012; Dantchev et al. 2007].

Dantchev et al. [2007] initiated the study of *parameterized proof complexity*. After considering the notions of propositional *parameterized tautologies* and *fpt-bounded* proof systems, they laid the foundations for the study of complexity of proofs in a parameterized setting. The problem BOUNDED CNF SAT leads to parameterized contradictions.

*Definition 2.1 (Dantchev et al. [2007]).* A *parameterized contradiction* is a pair  $(F, k)$  consisting of a propositional formula  $F$  in CNF and  $k \in \mathbb{N}$  such that  $F$  has no satisfying assignment of weight  $\leq k$ .

Here we focus on parameterized contradictions  $(F, k)$  where  $F$  is already unsatisfiable. We also study proof complexity for showing that a graph has no clique of size  $k$ , when in fact it does not contain any. The latter is an example in which the parameterization is implicit in the formula  $F$  itself. It is natural to also consider the case of parameterized contradictions  $(F, k)$  where  $F$  is satisfiable by assignments of weight greater  $k$ ; this case is ignored here, but some observations in this direction are made in Beyersdorff et al. [2012].

The notions of a parameterized proof system and of fpt-bounded proof systems were also developed by Dantchev et al. [2011].

*Definition 2.2 (Dantchev et al. [2011]).* A *parameterized proof system* for a parameterized language  $L \subseteq \Sigma^* \times \mathbb{N}$  is a function  $P : \Sigma^* \times \mathbb{N} \rightarrow \Sigma^* \times \mathbb{N}$  such that  $\text{rng}(P) = L$  and  $P(x, k)$  can be computed in time  $O(f(k)|x|^{O(1)})$  for some computable function  $f$ . The system  $P$  is *fpt-bounded* if there exist computable functions  $s$  and  $t$  such for that every  $(x, k) \in L$  there is  $(y, k')$  with  $P(y, k') = (x, k)$ ,  $|y| \leq s(k)|x|^{O(1)}$  and  $k' \leq t(k)$ .

A proof can be padded in order to ensure polynomial runtime, thus the original definition of parameterized proof systems can be simplified as follows.

*Definition 2.3 (alternative to Definition 2.2 suggested in [Beyersdorff et al. 2012]).* A *parameterized proof system* for a parameterized language  $L \subseteq \Sigma^* \times \mathbb{N}$  is a polynomial-time computable function  $P : \Sigma^* \rightarrow \Sigma^* \times \mathbb{N}$  such that  $\text{rng}(P) = L$ . The system  $P$  is *fpt-bounded* if there exists a computable function  $f$  such that for every  $(x, k) \in L$  there exists a proof  $y$  of size at most  $f(k)|x|^{O(1)}$  with  $f(y) = (x, k)$ .

<sup>2</sup>Compare BOUNDED CNF SAT to the canonical WEIGHTED CNF SAT problem, which asks for assignments of weight *exactly*  $k$ . The latter is also  $W[2]$ -complete [Flum and Grohe 2006] and is more common in the parameterized complexity literature.

The main motivation behind the work of Dantchev et al. [2011] was that of generalizing the classical approach of Cook and Reckhow [1979] to the parameterized case and working towards a separation of parameterized complexity classes by techniques developed in proof complexity.<sup>3</sup>

## 2.1. Parameterized Resolution and Parameterized DPLL

A *literal* is a positive or negated propositional variable and a *clause* is a set of literals, interpreted as their disjunction. The *width* of a clause is the number of its literals. A formula in conjunctive normal form (CNF) is a set of clauses, interpreted as their conjunction. The *Resolution system* is a refutation system for the set of all unsatisfiable CNF. Resolution gets its name from its only rule, the *Resolution rule*

$$\frac{\{x\} \cup C \quad \{\neg x\} \cup D}{C \cup D}$$

for clauses  $C, D$  and a variable  $x$ . The aim in Resolution is to demonstrate unsatisfiability of a clause set by deriving the empty clause from them. A Resolution refutation can be associated with directed acyclic graph where clauses in the refutation correspond to nodes and each inference  $\frac{C \quad D}{E}$  gives rise to two edges  $(C, E)$  and  $(D, E)$ . If this graph associated with the refutation is a tree, then the refutation is called *tree-like*, otherwise it is called *dag-like*. In a tree-like refutation any noninitial clause that is needed more than once must be rederived from scratch. The *size* of a Resolution proof is the number of its clauses where multiple occurrences of the same clause are counted separately. Undoubtedly, Resolution is the most studied and best-understood propositional proof system (cf. Segerlind [2007]).

For the remaining part of this article we will concentrate on *Parameterized Resolution* as introduced by Dantchev et al. [2007]. Parameterized Resolution is a refutation system for the set of parameterized contradictions (cf. Definition 2.1). Given a set of clauses  $F$  in variables  $x_1, \dots, x_n$ , a *Parameterized Resolution refutation* of  $(F, k)$  is a Resolution refutation of

$$F \cup \{\neg x_{i_1} \vee \dots \vee \neg x_{i_{k+1}} \mid 1 \leq i_1 < \dots < i_{k+1} \leq n\} .$$

Thus, in Parameterized Resolution we have built-in access to all parameterized clauses of the form  $\neg x_{i_1} \vee \dots \vee \neg x_{i_{k+1}}$ . All these clauses are available in the system, but when measuring the size of a refutation we only count those that occur in the refutation.

If refutations are tree-like we speak of *tree-like Parameterized Resolution*. Running parameterized DPLL procedures on parameterized contradictions produces tree-like Parameterized Resolution refutations, thus tree-like Resolution proof lengths are connected with the running time of DPLL procedures. Exactly as in usual tree-like Resolution, a tree-like Parameterized refutation of  $(F, k)$  can equivalently be described as a *boolean decision tree* where inner nodes are labeled with variables from  $F$  and leaves

<sup>3</sup>In fact, there are several Cook-type programs that can be associated with showing lower bounds for parameterized proof systems. The first paper of Dantchev et al. [2007] targets at FPT vs. W[2]. By giving a more general version of Definition 2.1 allowing arbitrary formulas  $F$  not necessarily in CNF, the journal version [Dantchev et al. 2011] recasts this as FPT vs. W[SAT]. In Beyersdorff et al. [2012] it is pointed out that a parameterized language  $L$  has an fpt-bounded proof system if and only if  $L$  is in para-NP, and hence we are aiming at a separation of coW[2] and para-NP (implying FPT  $\neq$  W[2]).

are labeled either with clauses from  $F$  or with parameterized clauses  $\neg x_{i_1} \vee \dots \vee \neg x_{i_{k+1}}$ . Each path in the tree corresponds to a partial assignment where a variable  $x$  gets value 0 or 1 according to whether the path branches left or right at the node labeled with  $x$ . The condition on the decision tree is that each path  $\alpha$  must lead to a clause that is falsified by the assignment corresponding to  $\alpha$ . Therefore, a boolean decision tree solves the *search problem* for  $(F, k)$ , which, given an assignment  $\alpha$ , asks for a clause falsified by  $\alpha$ . It is easy to verify that each tree-like Parameterized Resolution refutation of  $(F, k)$  yields a boolean decision tree for  $(F, k)$  and vice versa, where the size of the Resolution proof equals the number of nodes in the decision tree [Ben-Sasson et al. 2004].

### 3. ASYMMETRIC PROVER-DELAYER GAMES FOR DPLL REFUTATIONS

The original Prover-Delayer game for tree-like Resolution has been developed by Pudlák and Impagliazzo [2000]. In the game, Prover queries a variable and Delayer either gives it a value or leaves the decision to Prover and receives *one* point. The number of Delayer's points at the end of the game bounds from below the height of the proof tree.

This game has also been studied by Esteban and Torán [2003], where it is proved that the *clause space complexity* of a formula in tree-like Resolution is two plus the largest number of points achievable by the Delayer. Then the lower bound for the proof length follows from the fact that a formula with clause space complexity  $s$  requires proof length at least  $2^s - 1$  [Esteban and Torán 2001]. This connection to clause space complexity limits the strength of the method, since there are formulas for which the above lower bound is not tight (e.g., the classical pigeonhole principle). This is so because the clause space complexity of a formula  $F$  is  $s$  if and only if any proof tree for  $F$  contains a complete binary tree of height  $s$ . The gap between the size of this minor and the size of the proof tree is exactly what the original game fails to analyze.

The game used here, in contrast, assigns points to the Delayer asymmetrically ( $\log_2 c_0$  and  $\log_2 c_1$ ) according to two functions  $c_0$  and  $c_1$  (s.t.  $c_0^{-1} + c_1^{-1} = 1$ ) that depend on the principle, the variable queried, and the current partial assignment. In fact, the original Prover-Delayer game of Pudlák and Impagliazzo [2000] is the case where  $c_0 = c_1 = 2$ .

Loosely speaking, we interpret the inverse of the score functions as a way to define a distribution on the choices made by the DPLL algorithm. Under this view the Delayer's score at each step is just the entropy of the bit encoding the corresponding choice. Since root-to-leaf paths are in bijection with leaves, this process induces a distribution on the leaves. Hence the entropy collected on the path is the entropy of the corresponding leaf choice. In this interpretation, the asymmetric Prover-Delayer game becomes a challenge between Prover, who wants to end the game giving up little entropy, and Delayer, who wants to get a lot of it. This means that the average score of the Delayer is a measure (actually a lower bound) of the number of leaves. In our setup the DPLL algorithm decides the Prover queries, and the score function defines the distribution on paths. The role of the Delayer corresponds to a conditioning on this distribution.

We now describe the details of the game; in the following we intend all logarithms to be in base 2. Let  $(F, k)$  be a parameterized contradiction where  $F$  is a set of clauses in  $n$  variables  $x_1, \dots, x_n$ . We define a Prover-Delayer game: Prover and Delayer build a (partial) assignment to  $x_1, \dots, x_n$ . The game is over as soon as the partial assignment falsifies either a clause from  $F$  or a parameterized clause  $\neg x_{i_1} \vee \dots \vee \neg x_{i_{k+1}}$  where  $1 \leq i_1 < \dots < i_{k+1} \leq n$ . The game proceeds in rounds. In each round, Prover suggests a variable  $x_i$ , and Delayer either chooses a value 0 or 1 for  $x_i$  or leaves the choice to the Prover. In this last case the Prover sets the value and the Delayer gets some points.

The number of points Delayer earns depends on the variable  $x_i$ , the assignment  $\alpha$  constructed so far in the game, and two functions  $c_0$  and  $c_1$ . More precisely, the number of points that Delayer will get is

$$\begin{aligned} & 0 && \text{if Delayer chooses the value,} \\ & \log c_0(x_i, \alpha) && \text{if Prover sets } x_i \text{ to 0, and} \\ & \log c_1(x_i, \alpha) && \text{if Prover sets } x_i \text{ to 1.} \end{aligned}$$

Moreover, the functions  $c_0$  and  $c_1$  are nonnegative and are chosen in such a way that for each variable  $x$  and assignment  $\alpha$

$$\frac{1}{c_0(x, \alpha)} + \frac{1}{c_1(x, \alpha)} = 1 \quad (1)$$

holds. We remark that (1) is not strictly necessary for all  $\alpha$  and  $x$ , but it must hold at least for those assignments  $\alpha$  and choices  $x$  of the Delayer that can actually occur in any game with the Delayer strategy. We call this game the  $(c_0, c_1)$ -game on  $(F, k)$ . The connection of this game to size of proofs in tree-like Parameterized Resolution is given by the next theorem.

**THEOREM 3.1** [BEYERSDORFF ET AL. 2010B, 2010A]. *Let  $(F, k)$  be a parameterized contradiction and let  $c_0$  and  $c_1$  be two functions satisfying (1) for all partial assignments  $\alpha$  to the variables of  $F$ . If  $(F, k)$  has a tree-like Parameterized Resolution refutation of size at most  $S$ , then for each  $(c_0, c_1)$ -game played on  $(F, k)$  there is a Prover strategy (possibly dependent on the Delayer) that gives the Delayer at most  $\log S$  points.*

**PROOF.** Let  $(F, k)$  be a parameterized contradiction using variables  $x_1, \dots, x_n$ . Choose any treelike Parameterized Resolution refutation of  $(F, k)$  of size  $S$  and interpret it as a boolean decision tree  $T$  for  $F$ .

The decision tree  $T$  completely specifies the query strategy for Prover: at the first step he will query the variable labeling the root of  $T$ . Whatever decision is made regarding the value of the queried variable, Prover moves to the root of the corresponding subtree and queries the variable that labels it. This process induces a root-to-leaf walk on  $T$ , and such walks are in bijection with the set of leaves.

To completely specify Prover's strategy we need to explain how Prover chooses the value of the queried variable in case Delayer asks him to. The Prover is deterministic, but for the sake of this proof he will choose the answers at random: on average the Delayer score will be low enough to allow an easy derandomization of Prover's choices. A game position is completely described by the partial assignment  $\alpha$  computed so far, and by the variable  $x \notin \text{dom}(\alpha)$  queried at that moment. If the Prover is asked to answer the query for  $x$ , the answer will be:

$$\begin{cases} 0 & \text{with probability } \frac{1}{c_0(x, \alpha)} \\ 1 & \text{with probability } \frac{1}{c_1(x, \alpha)}. \end{cases}$$

Thus we are dealing with a randomized Prover strategy. In a game played between our randomized Prover and a specific Delayer  $D$ , we denote by  $p_{D, \ell}$  the probability of such a game to end at a leaf  $\ell$ . We call  $\pi_D$  this distribution on the leaves. To prove the theorem the following observation is crucial.

If the game ends at leaf  $\ell$ , then Delayer  $D$  scores exactly  $\log \frac{1}{p_{D, \ell}}$  points.

Before proving this claim, we show that it implies the theorem. The expected score of a Delayer  $D$  is

$$H(\pi_D) = \sum_{\ell} p_{D,\ell} \log \frac{1}{p_{D,\ell}},$$

which is the information-theoretic entropy of  $\pi_D$ . Since the support of  $\pi_D$  has size at most  $S$ , we obtain  $H(\pi_D) \leq \log S$ , because the entropy is maximized by the uniform distribution. By fixing the random choices of the Prover, we can force Delayer  $D$  to score at most  $\log S$  points.

To prove the claim consider a leaf  $\ell$  and the unique path that reaches it. W.l.o.g. we assume that this path corresponds to the ordered sequence of assignments  $x_1 = \epsilon_1, \dots, x_m = \epsilon_m$ . The probability of reaching the leaf is

$$p_{D,\ell} = p_1 p_2 \cdots p_m,$$

where  $p_i$  is the probability of setting  $x_i = \epsilon_i$  conditioned on the previous choices. If Prover chooses the value of the variable  $x_i$ , the score Delayer  $D$  gets at step  $i$  is

$$\log c_{\epsilon_i}(x_i, \{x_1 = \epsilon_1, x_2 = \epsilon_2, \dots, x_{i-1} = \epsilon_{i-1}\}),$$

which is exactly  $\log \frac{1}{p_i}$ . If Delayer makes the choice at step  $i$ , then  $p_i = 1$  and the score is 0, which is also  $\log \frac{1}{p_i}$ . Thus the score of the game play is

$$\sum_{i=1}^m \log \frac{1}{p_i} = \log \frac{1}{\prod_{i=1}^m p_i} = \log \frac{1}{p_{D,\ell}},$$

and this concludes the proof of the claim and the theorem.  $\square$

Notice that by setting  $c_0(x, \alpha) = c_1(x, \alpha) = 2$  for all variables  $x$  and partial assignments  $\alpha$ , we get the game of Pudlák and Impagliazzo [2000]. Proving lower bounds in our new game, that is, devising good Delayer strategies, entails first of all to finding suitable functions  $c_0$  and  $c_1$ . Functions  $c_0$  and  $c_1$  can be interpreted in terms of *information content* of tree-like Resolution refutations. The points that Delayer scores in one round should be proportional to the fraction of the current refutation that Prover can avoid to check by deciding a value for the variable. This is easily understandable in the case of the original game: the only good strategy for Prover to set a variable is choosing the value that allows him to proceed the game in the smallest of the two subtrees of the current refutation, which is in fact of size smaller than  $1/2$  of the current refutation size.

#### 4. TWO APPLICATIONS OF THE LOWER BOUND METHOD

In this section we show that parameterized contradictions based on the *pigeonhole principle* and on the *ordering principle* have lower bounds  $n^{k(1-o(1))}$  and  $n^{\sqrt{k}(1-o(1))}$ , respectively.<sup>4</sup> The main result of Dantchev et al. [2011] already implies a weaker lower bound of  $n^{k^{\Omega(1)}}$  for both formulas, and our results improve on that. It is interesting to remark that an  $n^{k/5}$  dag-like Parameterized Resolution lower bound for the pigeonhole principle is known from Beyersdorff et al. [2012].

<sup>4</sup>Although in the parameterized setting size functions have two arguments ( $n$  and  $k$ ), the notation  $o(1)$  refers to a function in  $n$  (alone).

We now focus on the *pigeonhole principle*  $PHP_n^{n+1}$ . Variable  $x_{i,j}$  for  $i \in [n+1]$  and  $j \in [n]$  indicates that pigeon  $i$  goes into hole  $j$ .  $PHP_n^{n+1}$  consists of the clauses

$$\bigvee_{j \in [n]} x_{i,j} \quad \text{for all pigeons } i \in [n+1]$$

and  $\neg x_{i_1,j} \vee \neg x_{i_2,j}$  for all choices of distinct pigeons  $i_1, i_2 \in [n+1]$  and holes  $j \in [n]$ . We prove that  $PHP_n^{n+1}$  is hard for tree-like Parameterized Resolution.<sup>5</sup>

**THEOREM 4.1.** *Any tree-like Parameterized Resolution refutation of  $(PHP_n^{n+1}, k)$  has size  $n^{k(1-o(1))}$ .*

**PROOF.** Let  $\alpha$  be a partial assignment to the variables  $\{x_{i,j} \mid i \in [n+1], j \in [n]\}$ . Let  $z_i(\alpha) = |\{j \in [n] \mid \alpha(x_{i,j}) = 0\}|$ , that is,  $z_i(\alpha)$  is the number of holes already excluded by  $\alpha$  for pigeon  $i$  (disregarding holes occupied by other pigeons). We define

$$c_0(x_{i,j}, \alpha) = \frac{n - z_i(\alpha)}{n - z_i(\alpha) - 1} \quad \text{and} \quad c_1(x_{i,j}, \alpha) = n - z_i(\alpha),$$

which clearly satisfies (1). We now describe Delayer's strategy in a  $(c_0, c_1)$ -game played on  $(PHP_n^{n+1}, k)$ . If Prover asks for a value of  $x_{i,j}$ , then Delayer decides as follows:

$$\begin{array}{ll} \text{set } \alpha(x_{i,j}) = 0 & \text{if there exists } i' \in [n+1] \setminus \{i\} \text{ such that } \alpha(x_{i',j}) = 1 \text{ or} \\ & \text{if there exists } j' \in [n] \setminus \{j\} \text{ such that } \alpha(x_{i,j'}) = 1 \\ \text{set } \alpha(x_{i,j}) = 1 & \text{if there is no } j' \in [n] \text{ with } \alpha(x_{i,j'}) = 1 \text{ and } z_i(\alpha) \geq n - k \\ \text{let Prover decide} & \text{otherwise.} \end{array}$$

Intuitively, Delayer leaves the choice to Prover as long as pigeon  $i$  does not already sit in a hole, there are at least  $k$  holes free for pigeon  $i$ , and there is no other pigeon sitting already in hole  $j$ . If Delayer uses this strategy, then clauses from  $PHP_n^{n+1}$  will not be violated in the game, that is, a contradiction will always be reached on some parameterized clause. To verify this claim, let  $\alpha$  be a partial assignment constructed during the game with  $w(\alpha) \leq k$  (we denote the weight of  $\alpha$  by  $w(\alpha)$ ). Then, for every pigeon that has not been assigned to a hole yet, there are at least  $k$  holes where it could go, and only  $w(\alpha)$  of these are already occupied by other pigeons. Thus  $\alpha$  can be extended to a one-one mapping of exactly  $k$  pigeons to holes.

Therefore, at the end of the game exactly  $k+1$  variables have been set to 1. Let us denote by  $p$  the number of variables set to 1 by Prover and let  $d$  be the number of 1's assigned by Delayer. As argued before  $p+d = k+1$ . Let us check how many points Delayer earns in this game. If Delayer assigns 1 to a variable  $x_{i,j}$ , then pigeon  $i$  was not assigned to a hole yet and, moreover, there must be  $n-k$  holes that are already excluded for pigeon  $i$  by  $\alpha$ , that is, for some  $J \subseteq [n]$  with  $|J| = n-k$  we have  $\alpha(x_{i,j'}) = 0$  for all  $j' \in J$ . Most of these 0's have been assigned by Prover, as Delayer has only assigned a 0 to  $x_{i,j}$  when some other pigeon was already sitting in hole  $j'$ , and there can be at most  $k$  such holes. Thus, before Delayer sets  $\alpha(x_{i,j}) = 1$ , she has already earned points for at least  $n-2k$  variables  $x_{i,j'}, j' \in J$ , yielding at least

$$\sum_{z=0}^{n-2k-1} \log \frac{n-z}{n-z-1} = \log \prod_{z=0}^{n-2k-1} \frac{n-z}{n-z-1} = \log \frac{n}{2k} = \log n - \log 2k$$

<sup>5</sup>In Theorem 4.1 and all following lower bounds it is implicitly understood that  $k$  is a fixed constant and  $n$  is sufficiently large with respect to  $k$ . This is the usual view adopted in parameterized complexity.

points for the Delayer. Note that because Delayer never allows a pigeon to go into more than one hole, she will earn at least the number of points calculated above for *each* of the  $d$  variables that she sets to 1.

If, conversely, Prover sets variable  $x_{i,j}$  to 1, then Delayer gets  $\log(n - z_i(\alpha))$  points for this, but she also receives points for most of the  $z_i(\alpha)$  variables set to 0 before that. Thus, in this case Delayer earns on pigeon  $i$  at least

$$\begin{aligned} & \log(n - z_i(\alpha)) + \sum_{z=0}^{z_i(\alpha)-k-1} \log \frac{n-z}{n-z-1} \\ &= \log(n - z_i(\alpha)) + \log \frac{n}{n - z_i(\alpha) + k} \\ &= \log n - \log \frac{n - z_i(\alpha) + k}{n - z_i(\alpha)} \\ &\geq \log n - \log k \end{aligned}$$

points. In total, Delayer gets at least

$$d(\log n - \log 2k) + p(\log n - \log k) \geq k(\log n - \log 2k)$$

points in the game. By Theorem 3.1, we obtain  $(\frac{n}{2k})^k$  as a lower bound to the size of each tree-like Parameterized Resolution refutation of  $(PHP_n^{n+1}, k)$ .  $\square$

In the above Delayer strategy we never associate a pigeon to more than one hole. For this reason the lower bound from Theorem 4.1 also holds for the *functional pigeonhole principle* where in addition to the clauses from  $PHP_n^{n+1}$  we also include  $\neg x_{i,j_1} \vee \neg x_{i,j_2}$  for all pigeons  $i \in [n+1]$  and distinct holes  $j_1, j_2 \in [n]$ . The Prover can easily force a Delayer to a partial assignment in which a hole does not receive any pigeon, so the proof does not work for the *bijective pigeonhole principle*, where all holes must host some pigeon.

As a second example we discuss the DPLL performance on the parameterized *ordering principle OP*, also called *least element principle*. The principle claims that any finite partially ordered set has a minimal element. There is a direct propositional translation of *OP* to a family  $OP_n$  of unsatisfiable CNFs. Each CNF  $OP_n$  expresses that there exists a partially ordered set of size  $n$  such that any element has a predecessor. The ordering principle has the following clauses.

$$\begin{array}{lll} \neg x_{i,j} \vee \neg x_{j,i} & \text{for every } i, j & \text{(Antisymmetry)} \\ \neg x_{i,j} \vee \neg x_{j,k} \vee x_{i,k} & \text{for every } i, j, k & \text{(Transitivity)} \\ \bigvee_{j \in [n] \setminus \{i\}} x_{j,i} & \text{for every } i & \text{(Predecessor)} \end{array}$$

With respect to parameterization the ordering principles are interesting. Both *OP* and the *linear ordering principle (LOP)*, which additionally assumes the order to be total, do not admit short tree-like Resolution refutations [Bonet and Galesi 2001] and have general Resolution refutations of polynomial size [Stalmark 1996]. In the parameterized setting things are different: *LOP* has short tree-like refutations [Beyersdorff et al. 2012] while *OP* does not and provides a separation between tree-like and dag-like Parameterized Resolution.

**THEOREM 4.2.** *Any tree-like Parameterized Resolution refutation of  $(OP_n, k)$  has size  $n^{\sqrt{k}(1-o(1))}$ .*

PROOF. Let  $\alpha$  be an assignment to the variables of  $OP$ . The Delayer will keep the following information:

- $G(\alpha) = (V(\alpha), E(\alpha))$  the graph obtained taking as edges the  $(i, j)$ 's such that  $\alpha(x_{i,j}) = 1$ ;
- $G^*(\alpha)$  the transitive closure of  $G(\alpha)$  and  $G^T(\alpha)$  the transpose graph of  $G(\alpha)$ .

In particular, for any vertex  $j$  in  $G(\alpha)$ , the Delayer considers the following information

- $z_j(\alpha) = |\{i \in [n] \mid \alpha(x_{i,j}) \text{ is not assigned}\}|$ ,
- $\dot{P}red_j(\alpha) = \{i \in [n] \mid \alpha(x_{i,j}) = 1\}$ , and
- $PPred_j(\alpha)$  the subset of  $\dot{P}red_j(\alpha)$  of those edges set to 1 by the Prover.

Loosely speaking the Delayer, taking as few decisions as possible, wants to force: (1) the game to end on a parameterized clause, and (2) the Prover to decide only one predecessor for each node. To reach the former, in some cases she will be forced to decide a predecessor of a node  $j$  to avoid that after few more trivial queries the game ends on a predecessor clause. To get (2) she will be forced to say that some node cannot be predecessor of some node  $j$ . In both cases we will prove that Delayer will keep her number of decisions bounded.

Let  $\alpha$  be the assignment built so far in the game and let  $x_{i,j}$  be the variable queried by Prover. Delayer acts as follows:

- (1) if  $(i, j) \in E(\alpha)^*$ , then she answers 1;
- (2) if  $(i, j) \in (E(\alpha)^*)^T$ , then she answers 0;
- (3) if  $|\dot{P}red_j(\alpha)| = 0$  and  $z_j(\alpha) \leq k + 1$ , then she answers 1;
- (4) if  $|PPred_j(\alpha)| \geq 1$ , then she answers 0;
- (5) otherwise, she leaves the decision to the Prover.

To simplify the argument we assume that in the game, after each decision by the Prover or after a decision by the Delayer according to Rule 3, the Prover asks all variables corresponding to edges that are in  $G^*(\alpha)$  and  $(G(\alpha)^*)^T$  but not in  $G(\alpha)$ . This will not change our result since on these nodes Delayer does not score any point.

Let  $P^\epsilon(t)$  be the set of edges set to  $\epsilon \in \{0, 1\}$  by the Prover after stage  $t$  ends. Let  $D^\epsilon(t)$  be the set of edges set to  $\epsilon \in \{0, 1\}$  by the Delayer. Finally, let  $D^*(t) \subseteq D^1(t)$  be the set of edges set to 1 by the Delayer according to Rule 3 of her strategy.  $P_j^\epsilon(t)$ ,  $D_j^\epsilon(t)$ , and  $D_j^*(t)$  are the subsets of the respective sets formed by those edges having end-node  $j$ , that is, edges of the form  $(i, j)$  for some  $i$ .

Let  $\alpha_t$  be the assignment built after stage  $t$  and let  $\alpha_t^*$  be the extensions of  $\alpha_t$  obtained by assigning all edges from  $G^*(\alpha_t)$  to 1 and all edges from  $(G(\alpha_t)^*)^T$  to 0. We define  $N_j(t) = \{(i, j) \mid i \in [n], (i, j) \in \text{dom}(\alpha_t^*) \setminus P^0(t)\}$ .

LEMMA 4.3. *At each stage  $t$  of the game, it holds that:*

- (1)  $|P^1(t)| + |D^*(t)| \geq \sqrt{|E(\alpha_t)|}$ ;
- (2) if  $w(\alpha_t) \leq k$  and  $|P_j^1(t)| + |D_j^*(t)| = 0$ , then  $|N_j(t)| \leq k$ ;
- (3) if  $w(\alpha_t) \leq k$ , then  $\alpha_t^*$  does not falsify any predecessor clause;
- (4) for each  $j \in [n]$ ,  $|D_j^*(t)| \leq 1$  and  $|P_j^1(t)| \leq 1$ .

PROOF. Condition 1 follows since  $|P^1(t)| + |D^1(t)| = |E(\alpha_t)|$ , and  $|E(\alpha_t)| \leq |E^*(\alpha_t)| \leq (|P^1(t)| + |D^*(t)|)^2$ .

Condition 2.  $|P_j^1(t)| + |D_j^*(t)| = 0$  implies that the vertex  $j$  has no predecessor. The only way to set a predecessor to a vertex that already has one is by Rule 1, but a

vertex without predecessors cannot get one by transitive closure. Then an edge  $x_{i,j}$  is in  $\text{dom}(\alpha_t^*) \setminus P^0(t)$  if and only if  $i$  is a successor of  $j$  in  $G^*(\alpha_t)$ . Hence there must be a directed tree rooted in  $j$  and containing all such successors. As  $\alpha_t$  has weight at most  $k$ ,  $G(\alpha_t)^T$  contains at most  $k$  edges and there are at most  $k$  successors of  $j$ . Therefore  $|N_j(t)| \leq k$ .

Condition 3. Consider a predecessor clause  $C_j$  that is not satisfied by  $\alpha_t$ . Then there are at least  $k + 1$  variables  $x_{i,j}$  unset, since otherwise, according to Rule 3 Delayer should have set one predecessor for  $j$ . If  $|P_j^1(t)| \geq 1$  or  $|D_j^*(t)| \geq 1$ , then  $C_j$  would be satisfied. Then by  $|P_j^1(t)| + |D_j^*(t)| = 0$  and by Condition 2 at most  $k$  additional literals of  $C_j$  are set to 0 by  $\alpha_t^*$ . The claim follows since there is at least one unset literal in  $C_j$ .

Condition 4. The first time that a predecessor of some node  $j$  is decided in the game is either by a decision of the Prover or by a decision of the Delayer according to Rule 3. Since Delayer applies Rule 3 only in the case no predecessor has been yet decided, it follows that  $|D_j^*(t)| \leq 1$ . Moreover, by Rule 4 Delayer prevents the Prover to set more than one predecessor for each node, hence  $|P_j^1(t)| \leq 1$ .  $\square$

LEMMA 4.4. *After the last stage  $f$  of the game the following holds:*

- a parameterized clause is falsified;
- $|P^1(f)| + |D^*(f)| \geq \sqrt{k + 1}$ .

PROOF. For the first condition, we notice that Rules 1 and 2 in the Delayer's strategy guarantee that neither antisymmetry nor transitivity axioms will be ever falsified during the game. Assuming that  $\alpha_f$  has weight strictly less than  $k + 1$ , then by Lemma 4.3 (condition 3), no predecessor clause is falsified. Hence  $w(\alpha_f) = k + 1$  and a parameterized clause is falsified.

The second property follows by Lemma 4.3 (condition 1) and by  $|E(\alpha_f)| \geq w(\alpha_f)$ , which is equal to  $k + 1$  because of the first part of this lemma.  $\square$

Set  $c_1(x_{i,j}, \alpha) = z_j(\alpha)$  and  $c_0(x_{i,j}, \alpha) = \frac{z_j(\alpha)}{z_j(\alpha) - 1}$ . For a given play of the game, let  $t_{i,j}$  be the stage of the game when the variable  $x_{i,j}$  is set. Let  $sc_j(t)$  be the number of points scored by the Delayer up to stage  $t$  for answers of the Prover to the variables  $x_{1,j}, x_{2,j}, \dots, x_{n,j}$ . Then the number of points scored by the Delayer at the end of the game is  $\sum_{j=1}^n sc_j(f)$ .

LEMMA 4.5. *The following implications hold*

- (1) *If  $|P_j^1(f)| = 1$ , then  $sc_j(f) \geq \log n - \log(k + 1)$ .*
- (2) *If  $|D_j^*(f)| = 1$ , then  $sc_j(f) \geq \log n - \log(2k + 1)$ .*

PROOF. For the first claim, let  $(i, j) \in D_j^*(f)$  and let  $t_{i,j}$  be the stage when  $x_{i,j}$  was set. We claim that  $|P_j^0(t_{i,j})| \geq n - (2k + 1)$ . W.l.o.g. we can assume that the variables  $x_{i',j}$  set to 0 by the Prover are the first ones with end-node  $j$  to be set to 0, because  $c_0(x_{i',j}, \alpha)$  is strictly decreasing with respect to  $z_j(\alpha)$ . Hence the Delayer gets at least

$$\sum_{l=2k+2}^n \log \frac{l}{l-1} = \log n - \log(2k + 1)$$

points on variables  $x_{1,j}, \dots, x_{n,j}$ .

It remains to prove the claim that  $|P_j^0(t_{i,j})| \geq n - (2k + 1)$ . According to Rule 3 of the strategy, there are at least  $n - (k + 1)$  variables  $x_{i',j}$  set to 0 in  $\alpha_{t_{i,j}}$ . Hence  $|P_j^0(t_{i,j})| + |D_j^0(t_{i,j})| \geq n - (k + 1)$ . Since at this stage  $i$  is the first predecessor of  $j$  to be fixed, then the Delayer has not set variables  $x_{i',j}$  to 0 according to Rule 4, but only by Rule 2.

Moreover, for the same reason, if  $t'$  is the stage preceding  $t_{i,j}$  we have that:  $|D_j^0(t_{i,j})| = |D_j^0(t')| = |N_j(t')| \leq k$ , where the last inequality holds by Lemma 4.3 (part 2). Then  $|P_j^0(t_{i,j})| \geq n - (2k + 1)$ .

We now show the second claim of the lemma. Let  $t_{i,j}$  be the stage in which Prover sets some  $x_{i,j}$  to 1, and let  $\alpha$  be the partial assignment corresponding to that stage. W.l. o. g. we assume that all variables in  $P_j^0(t_{i,j})$  are set before any variable in  $D_j^0(t_{i,j})$ , because  $c_0$  is monotone decreasing in the size of the second argument. Fix  $p = |P_j^0(t_{i,j})|$ . By Lemma 4.3 (part 2) we get  $|N_j(t')| \leq k$  where  $t'$  is the stage preceding  $t_{i,j}$ . Hence we know that  $z_j(\alpha) \geq n - k - p$ . The amount of points got by Delayer on vertex  $j$  is at least

$$\sum_{l=n-p+1}^n \log \frac{l}{l-1} + \log(n - k - p) = \log n - \log \frac{n-p}{n-k-p} \geq \log n - \log(k+1) .$$

□

The Delayer scores  $\sum_{j=1}^n sc_j(f)$ . By Lemma 4.4 there are at least  $\sqrt{k+1}$  vertices such that either  $|D_j^*(f)| \geq 1$  or  $|P_j^1(f)| \geq 1$ . For each vertex such events are mutually exclusive by the definition of the rules. Then by Lemma 4.5 Delayer gets at least  $\sqrt{k+1}(\log n - \log(2k+1))$  points. By Theorem 3.1 we get the lower bound.

## 5. TREELIKE RESOLUTION COMPLEXITY OF $K$ -CLIQUE

Instead of adding parameterized clauses of the form  $\neg x_{i_1} \vee \dots \vee \neg x_{i_{k+1}}$ , there are also more succinct ways to enforce only satisfying assignments of weight  $\leq k$ . One such method was considered by Dantchev et al. [2011] where for a formula  $F$  in  $n$  variables  $x_1, \dots, x_n$  and a parameter  $k$ , a new formula  $M = M(F, k)$  is computed such that  $F \wedge M$  is satisfiable if and only if  $F$  has a satisfying assignment of weight at most  $k$ . The formula  $M$  uses new variables  $s_{i,j}$ , where  $i \in [k]$  and  $j \in [n]$ , and consists of the clauses

$$\neg x_j \vee \bigvee_{i=1}^k s_{i,j} \quad \text{and} \quad \neg s_{i,j} \vee x_j \quad \text{for } i \in [k] \text{ and } j \in [n] \quad (2)$$

$$\neg s_{i,j} \vee \neg s_{i,j'} \quad \text{for } i \in [k] \text{ and } j \neq j' \in [n] \quad (3)$$

$$\neg s_{i,j} \vee \neg s_{i',j} \quad \text{for } i \neq i' \in [k] \text{ and } j \in [n]. \quad (4)$$

The clauses (2) express the fact that an index  $i$  is associated to a variable  $x_j$  if and only if this variable is set to true. The fact that the association is an injective function is expressed by the clauses (3) and (4).

This formula transformation not only leads to shorter formulas; it also allows smaller parameterized refutations in some cases (e.g., for the pigeonhole principle in dag-like Resolution [Beyersdorff et al. 2012; Dantchev et al. 2011]).

Beyersdorff et al. [2012] argue that the clique formula is “invariant” with respect to this transformation, thus its classical proof complexity is equivalent to its parameterized proof complexity (in both the formulation with explicit parameterized axioms and the succinct encoding). Therefore Beyersdorff et al. [2012] posed the question of

determining the complexity of the clique formulas in classical Resolution. Theorem 5.2 below provides an answer to this question for the tree-like case.

### 5.1. Random $k$ -colorable Graphs

Our study focuses on the average-case complexity of proving the absence of a  $k$ -clique in random graphs distributed according to a variation of the Erdős-Rényi model  $G(n, p)$ . In this Erdős-Rényi model, random graphs on  $n$  vertices are constructed by including every edge independently with probability  $p$ . It is known that  $k$ -cliques appear at the threshold probability  $p^* = n^{-\frac{2}{k-1}}$ . If  $p < p^*$ , then with high probability there is no  $k$ -clique; while for  $p > p^*$  with high probability there are many. For  $p = p^*$  there is a  $k$ -clique with constant probability.

The complexity of  $k$ -clique has already been studied in restricted computational models by Rossman [2008, 2010]. He shows that in these models any circuit that succeeds with good probability on graph distributions close to the critical threshold requires size  $\Omega(n^{\frac{k}{4}})$ , and even matching upper bounds exist in these models [Amano 2010; Rossman 2010]. Since we want to study negative instances of the clique problem, we focus on probability distributions with  $p < p^*$ . To simplify the proof we will prove a lower bound for a slightly sparser distribution. We now give the CNF formulation of a statement claiming that a  $k$ -clique exists in a graph.

*Definition 5.1.* Given a graph  $G = (V, E)$  and a parameter  $k$ ,  $Clique(G, k)$  is a formula in conjunctive normal form containing the following clauses

$$\bigvee_{v \in V} x_{i,v} \quad \text{for every } i \in [k] \quad (5)$$

$$\neg x_{i,u} \vee \neg x_{j,v} \quad \text{for every } i, j \in [k], i \neq j \text{ and every } \{u, v\} \notin E \quad (6)$$

$$\neg x_{i,u} \vee \neg x_{i,v} \quad \text{for every } u \neq v \in V. \quad (7)$$

The interpretation of variable  $x_{i,v}$  is that vertex  $v$  is the  $i$ th member of the clique. Clearly the formula  $Clique(G, k)$  is satisfiable if and only if the graph  $G$  has a clique of size  $k$ .

We now describe a family of hard graph instances for  $k$ -clique. We consider a random graph  $G$  on  $kn$  vertices. The set of vertices  $V$  is divided into  $k$  blocks of  $n$  vertices each, named  $V_1, V_2, \dots, V_k$ . Edges may be present only between vertices of different blocks, and these edges are chosen independently at random. For any constant  $\epsilon$  and any pair of vertices  $(u, v)$  with  $u \in V_i, v \in V_{i'}$  and  $i < i'$ , the edge  $\{u, v\}$  is present with probability

$$p = n^{-(1+\epsilon)\frac{2}{k-1}}.$$

We call this distribution of graphs  $\mathcal{G}_{k,\epsilon}$ . Notice that all graphs in  $\mathcal{G}_{k,\epsilon}$  are properly colorable with  $k$  colors. Later we will focus on a specific range for  $\epsilon$ .

In a  $k$ -colorable graph  $G$  with color classes  $V_1, \dots, V_k$  a  $k$ -clique contains exactly one vertex per color class. In this case we can simplify formula  $Clique(G, k)$  by setting  $x_{i,v} = 0$  for every  $i \in [k]$  and  $v \in V_j$  such that  $i \neq j$ . Essentially we are forcing the  $i$ th vertex in the clique to be in the  $i$ th block. If variable  $x_{i,v}$  survives to the restriction then  $v \in V_i$ , thus the index  $i$  is redundant and can be dropped. The resulting simplified formula is the following:

$$\bigvee_{v \in V_i} x_v \quad \text{for every } i \in [k] \quad (8)$$

$$\neg x_u \vee \neg x_v \quad \text{for every } \{u, v\} \notin E(G). \quad (9)$$

Variable restrictions only reduce the size of a (tree-like) Resolution refutation. Therefore, in order to prove lower bounds for  $\text{Clique}(G, k)$  on graphs distributed according to  $\mathcal{G}_{k, \epsilon}$  we can focus on the simpler CNF formulas (8)–(9).

A comment regarding the encoding is required. Beame et al. [2007] studied formulas similar to  $\text{Clique}(G, k)$  for the dual problem of independent sets in the context of classical nonparameterized proof complexity. They study the case of  $k = \Omega(n)$ , so the former encoding has an interesting lower bound because it contains clauses of a nontrivial pigeonhole principle. In the parameterized framework this is not necessarily true, since  $k$  is small and  $\text{PHP}_{k-1}^k$  is considered feasible here.

We will now show that for a random graph  $G$  distributed according to  $\mathcal{G}_{k, \epsilon}$ , it holds with high probability that the smallest decision tree that proves unsatisfiability of  $\text{Clique}(G, k)$  has size  $n^{\Omega(k(1-\epsilon))}$ . To show that  $\text{Clique}(G, k)$  requires refutations of size  $n^{\Omega(k(1-\epsilon))}$  it suffices to exhibit two score functions  $c_0$  and  $c_1$  and a Delayer strategy such that the Delayer is guaranteed to score  $\Omega(k(1-\epsilon) \log n)$  points in any game played against any Prover.

**THEOREM 5.2.** *Let  $\epsilon$  be a constant such that  $0 < \epsilon < 1$ . For a random graph  $G$  distributed according to  $\mathcal{G}_{k, \epsilon}$ , it holds with high probability that the smallest treelike Parameterized Resolution refutation of  $\text{Clique}(G, k)$  has size  $n^{\Omega(k(1-\epsilon))}$ .*

**PROOF.** Let  $G$  be a random graph distributed according to  $\mathcal{G}_{k, \epsilon}$ . We prove a lower bound for the restricted formula (8)–(9): this immediately implies the statement of the theorem.

For any vertex set  $S$ , let  $\Gamma^c(S)$  be the set of vertices that are common neighbors of  $S$  (i. e., vertices that are connected to all vertices in  $S$ ). Notice that  $\Gamma^c(\emptyset)$  is  $V(G)$ . We first show that with high probability the following properties hold:

- (1)  $G$  has no clique of size  $k$ ;
- (2) For any set  $S$  of less than  $\frac{k}{4}$  vertices in distinct blocks,  $|\Gamma^c(S) \cap V_b| \geq n^{\Omega(1-\epsilon)}$  for any block  $V_b$  disjoint from  $S$ .

For item 1. The expected number of  $k$ -cliques in  $G$  is  $n^k p^{\binom{k}{2}} = n^{-k\epsilon}$ . By Markov inequality, the probability of the existence of a single  $k$ -clique is at most the expected value.

For item 2. It is sufficient to show the statement for sets of size exactly  $\frac{k}{4} - 1$ . Fix any such set  $S$ , and fix any block  $V_b$  that does not contain vertices in this set. We denote by  $X_i$  the random variable that is 1 when  $i \in \Gamma^c(S)$ , and 0 otherwise. Thus the size of  $V_b \cap \Gamma^c(S)$  is the sum of  $n$  independent variables. Notice that  $X_i$  is 1 with probability  $p^{\frac{k}{4}-1} \geq n^{-\frac{1+\epsilon}{2}}$ . Thus the expected value is at least  $n^{\frac{1-\epsilon}{2}}$ . We define

$$T = \frac{n^{\frac{1-\epsilon}{2}}}{2}.$$

Since  $T = n^{\Omega(1-\epsilon)}$  and  $T$  is a constant fraction of the expected value, by the Chernoff bound (see, for example, [Dubhashi and Panconesi 2009, Theorem 1.1]) we obtain that  $V_b \cap \Gamma^c(S)$  has size less than  $T$  with probability at most  $e^{-n^{\Omega(1-\epsilon)}}$ . By the union bound on the choices of block  $V_b$  and of set  $S$  of size  $\frac{k}{4} - 1$  we get item 2.

Thus we can conclude that with high probability the random graph  $G$  fulfills both 1 and 2.

We now define functions  $c_0$  and  $c_1$ , which are legal cost functions for an asymmetric Prover-Delayer game played on the  $k$ -clique formula of the graph  $G$ . We also exhibit

a Delayer strategy that is guaranteed to score  $\Omega(k \log T)$  points. This, together with Theorem 3.1, implies the main statement.

For any partial assignment  $\alpha$  we consider the set of vertices “chosen by  $\alpha$ ,” which is  $\{u \mid \alpha(x_u) = 1\}$ ; any vertex that is the common neighbor of the chosen set is called “good for  $\alpha$ .” Notice that if  $v$  is a good vertex for  $\alpha$ , then the partial assignment  $\alpha \cup \{x_v = 1\}$  does not falsify any clause in the formula. Notice also that  $\alpha$  may set to 0 some good vertices. In particular we denote by  $R_b(\alpha)$  the vertices of the block  $V_b$  that are good for  $\alpha$ , but are nevertheless set to 0 in  $\alpha$ :

$$R_b(\alpha) = \{v \in V_b \mid \alpha(x_v) = 0 \text{ and for all } u, \alpha(x_u) = 1 \text{ implies } \{u, v\} \in E(G)\}.$$

When asked for an unassigned variable  $x_v$ , for some  $v \in V_b$ , the Delayer applies the following rules in the given order:

- (1) If  $\alpha$  contains at least  $\frac{k}{4}$  variables set to 1, the Delayer surrenders;
- (2) if  $\alpha(x_u) = 1$  for some  $u$  with  $\{u, v\} \notin E(G)$  then Delayer answers 0;

(notice that at this point vertex  $v$  must be good for  $\alpha$ )

- (1) if  $R_b(\alpha)$  has size at least  $T - 1$ , then the Delayer answers 1;
- (2) otherwise the Delayer leaves the answer to the Prover.

During the game the invariant  $|R_b(\alpha)| < T$  holds for every  $b \in [k]$ : the only way such a set can increase in size is when Prover sets a good vertex in  $V_b$  to 0. Thus the size of  $R_b(\alpha)$  can only increase one by one. When it reaches  $T - 1$  and the Delayer is asked for a variable in that block, she will reply 1, so the size of  $R_b(\alpha)$  won't increase any more.

Another important property of the Delayer strategy is that her decision to answer 1 never falsifies a clause, since all blocks contain at least  $T$  good vertices at any moment during the game: as soon some  $R_b(\alpha)$  reaches size  $T$  the Delayer sets a vertex for block  $b$ . This is possible because of item 5.1 together with the fact that at most  $\frac{k}{4}$  vertices are set in  $\alpha$ . This proves that no clause in (8) can be falsified during the game.

Neither clauses in (9) can be falsified during the game: the Delayer imposes answer 0 whenever a vertex is not good for  $\alpha$ , which means that, if chosen, it would not form a clique with the ones chosen before. It is also not possible that the game ends by violating a parameterized clause as these are just weakenings of the clauses (9). Therefore, the game only ends when the Delayer gives up.

For an assignment  $\alpha$  and a vertex  $v \in V_b$ , let

$$c_0 = \frac{T - |R_b(\alpha)|}{T - |R_b(\alpha)| - 1} \quad \text{and} \quad c_1 = T - |R_b(\alpha)|.$$

Because of the previous observations the values of  $c_0$  and  $c_1$  are always nonnegative. Furthermore notice that when  $|R_b(\alpha)| = T - 1$  Delayer never leaves the choice to Prover, thus  $c_0$  is always well defined when the Delayer scores.

Consider a game play and the set of  $\frac{k}{4}$  vertices chosen by the final partial assignment  $\alpha$ . We show that for any chosen vertex, the Delayer scores  $\log T$  points for queries in the corresponding block.

Fix the block  $b$  of a chosen vertex  $u$ . Consider the assignment  $\alpha$  that corresponds to the game step when  $x_u$  is set to 1. Consider  $R = R_b(\alpha)$ . We identify partial assignments

$$\alpha_0 \subset \alpha_1 \subset \dots \subset \alpha_{|R|-1} \subset \alpha$$

corresponding to the moments in the game when Prover sets to 0 one of the variables indexed by  $R$ . For such rounds the Delayer gets at least

$$\sum_{i=0}^{|R|-1} \log \frac{T - |R_b(\alpha_i)|}{T - |R_b(\alpha_i)| - 1} \geq \sum_{i=0}^{|R|-1} \log \frac{T - i}{T - i - 1} = \log(T) - \log(T - |R|)$$

points. Here the first inequality follows from the fact that any vertex that is good at some stage of the game is also good in all previous stages. Thus  $|R_b(\alpha_i)| \geq i$ .

Now we must consider two cases: either  $x_u = 1$  is set by Prover, or it is set by Delayer. In the former case Delayer gets  $\log(T - |R|)$  points for Prover setting  $x_u = 1$ . Together with the points for the previous zeros this yields  $\log T$  points. In the latter case Delayer gets 0 points as she set  $x_u = 1$  by herself, but now  $|R| = T - 1$  and she got already  $\log T$  points for all the zeros assigned by Prover. In both cases the total score of the Delayer is  $\log T = \frac{1-\epsilon}{2} \log n$ .

Since this score is obtained in at least  $\frac{k}{4}$  blocks, we are done.  $\square$

## 5.2. Complete $(k - 1)$ -partite Graphs

Instead of random graphs, we are now looking at one of the canonical graphs without a  $k$ -clique: the  $(k - 1)$ -partite graph. We will show in this subsection that the  $k$ -clique formulas on complete  $(k - 1)$ -partite graphs has a “short” dag-like Resolution refutations. In contrast, we will show a lower bound for tree-like Resolution. Let  $C_n$  be the complete  $(k - 1)$ -partite graph in which each partition has size  $n$ . The formula  $\text{Clique}(C_n, k)$  claims that there is a way to place  $k$  indexes on the graph in such a way that no two indexes fall into the same partition. That essentially implies an injective mapping from  $k$  to  $k - 1$ , so it is unsatisfiable.

**PROPOSITION 5.3.** *The formulas  $\text{Clique}(C_n, k)$  have  $O(2^k k^2 n^2)$  Resolution refutations.*

**PROOF.** We follow the idea of a *monotone Resolution* (cf. [Buss and Pitassi 1998]) refutation of  $\text{PHP}_{k-1}^k$ . A monotone refutation of the pigeonhole principle in variables  $p_{i,j}$  contains disjunctions of positive literals as lines in the proof. The only available rule is

$$\frac{A \vee \bigvee_{i \in I_0} p_{i,h} \quad B \vee \bigvee_{i \in I_1} p_{i,h}}{A \vee B \vee \bigvee_{i \in I_0 \cap I_1} p_{i,h}}, \quad (10)$$

where  $h \in [k - 1]$  is a hole and  $I_0, I_1 \subseteq [k]$  are sets of pigeons. Clearly, the formulas  $\text{PHP}_{k-1}^k$  admit monotone Resolution refutations depending in size only on  $k$ .

For  $1 \leq h < k - 1$ , let  $V_h$  be the  $h$ th block in the partition of the vertices. We apply the following substitution in the monotone proof of  $\text{PHP}_{k-1}^k$

$$p_{i,h} \longleftrightarrow \bigvee_{v \in V_h} x_{i,v}.$$

To simulate an application of the inference rule (10) it is sufficient to show the inference for empty  $A$  and  $B$  and for  $I_0 \cap I_1 = \emptyset$ , since the simulation in its full generality then follows by weakening. Given  $\bigvee_{i \in I_0} \bigvee_{v \in V_h} x_{i,v}$  and  $\bigvee_{i \in I_1} \bigvee_{v \in V_h} x_{i,v}$  we know that for each  $i_0 \in I_0, i_1 \in I_1, v, w \in V_h$ , the clause  $\neg x_{i_0,v} \vee \neg x_{i_1,w}$  is an axiom of  $\text{Clique}(C_n, k)$ . Thus by using  $O(|V_h|^2 \cdot |I_0| \cdot |I_1|)$  of these axioms we easily get the empty clause.

There are monotone  $\text{PHP}_{k-1}^k$  refutations of size  $O(2^k)$ . Therefore the whole  $\text{Clique}(C_n, k)$  refutation has size at most  $O(2^k k^2 n^2)$ .  $\square$

An easy analysis of the above proof shows that  $\text{Clique}(C_n, k)$  even allows short *regular* Resolution refutations of size  $f(k)n^2$  for some function  $f$ . Regular Resolution is defined by the property that on each path from the root to a leaf each variable can be resolved at most once. Regular Resolution is an interesting subsystem of Resolution that is known to be more powerful than tree-like Resolution [Bonet and Galesi 2001], but weaker than full dag-like Resolution [Alekhovich et al. 2007].

The proof in Theorem 5.3 is based on the simulation of monotone Resolution by general (dag-like) Resolution [Buss and Pitassi 1998]. In general, this simulation is not possible for tree-like Resolution. In particular, in the previous proof our simulation of one monotone inference step is not tree-like. Hence, even if we started with a tree-like proof of  $\text{PHP}_{k-1}^k$ , the resulting proof of  $\text{Clique}(C_n, k)$  would be dag-like. Indeed, we show a lower bound for the clique formulas on  $(k - 1)$ -partite graphs for tree-like Resolution.

**THEOREM 5.4.** *Any tree-like Resolution refutation of  $\text{Clique}(C_n, k)$  requires size  $n^{\Omega(k)}$ .*

**PROOF.** The set of vertices of the graph  $C_n$  is partitioned into the sets  $V_1, \dots, V_{k-1}$  of size  $n$  each. We define a Delayer strategy such that at the end of the game the partial assignment always has  $k - 1$  indexes assigned to specific vertices in different blocks. Moreover, we will define the score functions in such a way that on each block Delayer scores exactly  $\log n$  points. That will conclude the proof.

Delayer keeps  $k - 1$  sets  $Z_1, \dots, Z_{k-1}$  with  $Z_j \subseteq V_j$  that represent the excluded vertices in each block. At the beginning of the game they are all empty. When the Prover has a partial assignment  $\alpha$  and queries  $x_{i,v}$  for  $v \in V_j$  the Delayer answers

- (1) 0 if  $x_{i,w}$  is true in  $\alpha$  for some  $w \neq v$ ;
- (2) 0 if  $x_{l,w}$  is true in  $\alpha$  for some  $l \in [k] \setminus \{i\}$  and some  $w \in V_j$ ;
- (3) 0 if  $v \in Z_j$ ;
- (4) 1 if  $v \notin Z_j$  and  $Z_j = V_j \setminus \{v\}$ ;
- (5) and lets the Prover decide otherwise.

After each round, Delayer updates the sets  $Z_j$  as follows. If Delayer sets the variable herself, then  $Z_j$  remains unaltered. Otherwise, if Prover decides 0 then Delayer sets  $Z_j := Z_j \cup \{v\}$ . If Prover decides 1, then  $Z_j := V_j \setminus \{v\}$ .

For his choices, Prover scores according to the functions

$$c_0 = \frac{|V_j| - |Z_j|}{|V_j| - |Z_j| - 1} \quad \text{and} \quad c_1 = |V_j| - |Z_j|.$$

Because of the first two rules of the Delayer strategy, the game always ends with a partial injective assignment of indexes to vertices in different blocks. Thus the only case in which the Delayer loses is when there is an index  $i$  such that  $\alpha$  sets  $x_{i,v} = 0$  for all vertices  $v$  in  $C_n$ . We claim that at the end of the game  $k - 1$  indexes have been assigned, one in each block. To prove this claim, assume for a contradiction that no index was assigned into the block  $V_j$ . Consider the last moment in the game in which  $x_{i,v} = 0$  has been assigned for some  $v \in V_j$ . By assumption, all variables  $x_{i,u}$  for  $u \in V_j \setminus \{v\}$  have been queried before and were already answered by 0. According to the Delayer strategy, either  $x_{i,u} = 0$  was set by Delayer by rule 3, or  $x_{i,u} = 0$  was decided by Prover. Hence in both cases  $u \in Z_j$  and therefore  $Z_j = V_j \setminus \{v\}$ . But then Delayer would assign  $x_{i,v}$  to 1 according to item 4 of her strategy, a contradiction.

Thus in the final configuration of the game  $k - 1$  indexes have been assigned. A counting argument similar (but simpler) to the ones in the previous proofs show that the Delayer scores exactly  $\log n$  points for the queries in each of the blocks.

Fix a block  $i$ : exactly one variable  $x_{i,v}$  is set to one. Let us say that  $|Z_i| = z$  right before that decision. Until that moment  $|Z_i|$  increases one by one every time Delayer scores some point on Prover deciding for some  $x_{i,u}$  to be zero. Delayer scores

$$\sum_{t=0}^{z-1} \log \frac{|V_i| - t}{|V_i| - t - 1} = \log |V_i| - \log(|V_i| - z).$$

Delayer chooses to set  $x_{i,v} = 1$  if and only if  $z = |V_i| - 1$ , otherwise the Prover chooses which gives  $\log(|V_i| - z)$  points to Delayer. In both cases Delayer scores  $\log |V_i|$  points on block  $i$ . Thus in the end, Delayer gets exactly  $(k - 1) \log n$  points.  $\square$

## 6. CONCLUSION AND OPEN PROBLEMS

In this article we have shown that the asymmetric Prover-Delayer game can be used to lower bound tree-like Resolution refutations. While the result itself was already present in a previous paper by the authors [Beyersdorff et al. 2010b], we provide here a new proof using a compelling information-theoretic argument. In order to show some example application of this game we strengthen the known lower bounds for tree-like parameterized Resolution for the pigeonhole principle and the ordering principle. A more interesting application is the  $n^{\Omega(k)}$  lower bound for  $\text{Clique}(G, k)$  formulas for tree-like Resolution.

It is natural to ask how hard it is to certify that a graph has no  $k$ -clique, compared to the brute force approach. A Resolution refutation of  $\text{Clique}(G, k)$  provides a sound witness of the absence of a  $k$ -clique, but how long is the refutation of  $\text{Clique}(G, k)$  in dag-like Resolution? A lower bound  $n^{\Omega(k)}$  for this formula would imply parameterized Resolution lower bounds even for the stronger parameterization of CNFs.

Another interesting fact about the  $k$ -clique formulas (as we have defined it) arises from their translation to 3-CNFs. These translated formulas have constant width and dag-like Resolution refutations of width  $\Theta(k)$  and size  $n^{O(k)}$ . A matching lower bound would indicate that clique formulas are the hardest among the ones of width  $\Theta(k)$ .

This is particularly interesting in light of results of Atserias et al. [2011], who prove that clause learning SAT solvers run in time  $n^{O(w)}$  if the minimal proof width is  $w$ . It is a natural question whether this is tight. If the clique formulas require proof size  $n^{\Omega(k)}$  this would show that the answer is yes.

## ACKNOWLEDGMENTS

We thank the anonymous referees of both the conference and the journal version of this article for their insightful suggestions that helped to improve it. We thank Pavel Pudlák for the second open problem.

## REFERENCES

- Alekhovich, M. and Razborov, A. A. 2008. Resolution is not automatizable unless W[P] is tractable. *SIAM J. Comput.* 38, 4, 1347–1363.
- Alekhovich, M., Johannsen, J., Pitassi, T., and Urquhart, A. 2007. An exponential separation between regular and general resolution. *Theory Comput.* 3, 1, 81–102.
- Amano, K. 2010. Subgraph isomorphism on  $\text{AC}^0$  circuits. *Comput. Complexity* 19, 2, 183–210.
- Atserias, A., Fichte, J. K., and Thurley, M. 2011. Clause-learning algorithms with many restarts and bounded-width resolution. *J. Artif. Intell. Res.* 40, 353–373.
- Beame, P. and Pitassi, T. 1996. Simplified and improved resolution lower bounds. In *Proceedings of the 37th IEEE Symposium on the Foundations of Computer Science*. 274–282.
- Beame, P., Karp, R. M., Pitassi, T., and Saks, M. E. 2002. The efficiency of resolution and Davis–Putnam procedures. *SIAM J. Comput.* 31, 4, 1048–1075.
- Beame, P., Kautz, H. A., and Sabharwal, A. 2004. Towards understanding and harnessing the potential of clause learning. *J. Artif. Intell. Res.* 22, 319–351.

- Beame, P., Impagliazzo, R., and Sabharwal, A. 2007. The resolution complexity of independent sets and vertex covers in random graphs. *Comput. Complex.* 16, 3, 245–297.
- Beame, P. W., Impagliazzo, R., Krajíček, J., Pitassi, T., and Pudlák, P. 1996. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. *Proc. London Math. Soc.* 73, 3, 1–26.
- Ben-Sasson, E. and Wigderson, A. 2001. Short proofs are narrow - resolution made simple. *J. ACM* 48, 2, 149–169.
- Ben-Sasson, E., Impagliazzo, R., and Wigderson, A. 2004. Near optimal separation of tree-like and general resolution. *Combinatorica* 24, 4, 585–603.
- Beyersdorff, O., Galesi, N., and Lauria, M. 2010a. Hardness of parameterized resolution. Tech. rep. TR10-059, Electronic Colloquium on Computational Complexity.
- Beyersdorff, O., Galesi, N., and Lauria, M. 2010b. A lower bound for the pigeonhole principle in treelike resolution by asymmetric prover-delayer games. *Inf. Process. Lett.* 110, 23, 1074–1077.
- Beyersdorff, O., Galesi, N., and Lauria, M. 2011. Parameterized complexity of DPLL search procedures. In *Proceedings 14th International Conference on Theory and Applications of Satisfiability Testing*. Lecture Notes in Computer Science, vol. 6695, Springer, 5–18.
- Beyersdorff, O., Galesi, N., Lauria, M., and Razborov, A. 2012. Parameterized bounded-depth Frege is not optimal. *ACM Trans. Comput. Theory* 4, 3.
- Blake, A. 1937. Canonical expressions in boolean algebra. Ph.D. thesis, University of Chicago.
- Bonet, M. L. and Galesi, N. 2001. Optimality of size-width tradeoffs for resolution. *Comput. Complexity* 10, 4, 261–276.
- Bonet, M. L., Esteban, J. L., Galesi, N., and Johannsen, J. 2000. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM J. Comput.* 30, 5, 1462–1484.
- Buss, S. R. and Pitassi, T. 1998. Resolution and the weak pigeonhole principle. In *Computer Science Logic*, Springer, 149–156.
- Chen, Y. and Flum, J. 2008. The parameterized complexity of maximality and minimality problems. *Ann. Pure Appl. Logic* 151, 1, 22–61.
- Chvátal, V. and Szemerédi, E. 1988. Many hard examples for resolution. *J. ACM* 35, 4, 759–768.
- Cook, S. A. and Reckhow, R. A. 1979. The relative efficiency of propositional proof systems. *J. Symb. Logic* 44, 1, 36–50.
- Dantchev, S. S., Martin, B., and Szeider, S. 2007. Parameterized proof complexity. In *Proceedings of the 48th IEEE Symposium on the Foundations of Computer Science*. 150–160.
- Dantchev, S. S., Martin, B., and Szeider, S. 2011. Parameterized proof complexity. *Comput. Complexity* 20, 51–85.
- Davis, M. and Putnam, H. 1960. A computing procedure for quantification theory. *J. ACM* 7, 210–215.
- Davis, M., Logemann, G., and Loveland, D. W. 1962. A machine program for theorem-proving. *Commun. ACM* 5, 7, 394–397.
- Downey, R. G. and Fellows, M. R. 1999. *Parameterized Complexity*. Springer.
- Dubhashi, D. P. and Panconesi, A. 2009. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press.
- Esteban, J. L. and Torán, J. 2001. Space bounds for resolution. *Inform. Comput.* 171, 1, 84–97.
- Esteban, J. L. and Torán, J. 2003. A combinatorial characterization of treelike resolution space. *Inf. Process. Lett.* 87, 6, 295–300.
- Flum, J. and Grohe, M. 2006. *Parameterized Complexity Theory*. Springer.
- Gao, Y. 2009. Data reductions, fixed parameter tractability, and random weighted d-CNF satisfiability. *Artif. Intell.* 173, 14, 1343–1366.
- Haken, A. 1985. The intractability of resolution. *Theor. Comput. Sci.* 39, 297–308.
- Janson, S., Łuczak, T., and Ruciński, A. 2000. *Random Graphs*. Wiley.
- Niedermeier, R. 2006. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics and Its Applications, Oxford University Press.
- Pudlák, P. and Impagliazzo, R. 2000. A lower bound for DLL algorithms for SAT. In *Proceedings of the 11th Symposium on Discrete Algorithms*. 128–136.
- Robinson, J. A. 1965. A machine-oriented logic based on the resolution principle. *J. ACM* 12, 23–41.
- Rossmann, B. 2008. On the constant-depth complexity of  $k$ -clique. In *Proceedings of the 40th ACM Symposium on Theory of Computing*. 721–730.
- Rossmann, B. 2010. The monotone complexity of  $k$ -clique on random graphs. In *Proceedings of the 51th IEEE Symposium on the Foundations of Computer Science*. IEEE, 193–201.

- Segerlind, N. 2007. The complexity of propositional proofs. *Bull. Symb. Logic* 13, 4, 417–481.
- Segerlind, N., Buss, S. R., and Impagliazzo, R. 2004. A switching lemma for small restrictions and lower bounds for  $k$ -DNF resolution. *SIAM J. Comput.* 33, 5, 1171–1200.
- Stalmark, G. 1996. Short resolution proofs for a sequence of tricky formulas. *Acta Informatica* 33, 277–280.
- Urquhart, A. 1987. Hard examples for resolution. *J. ACM* 34, 1, 209–219.

Received November 2011; revised August 2012; accepted October 2012