Size, Cost, and Capacity: A Semantic Technique for Hard Random QBFs*

Olaf Beyersdorff[1], Joshua Blinkhorn[2], and Luke Hinde[3]

1 **School of Computing, University of Leeds, UK**
`o.beyersdorff@leeds.ac.uk`
2 **School of Computing, University of Leeds, UK**
`scjlb@leeds.ac.uk`
3 **School of Computing, University of Leeds, UK**
`sclpeh@leeds.ac.uk`

Abstract

As a natural extension of the SAT problem, an array of proof systems for quantified Boolean formulas (QBF) have been proposed, many of which extend a propositional proof system to handle universal quantification. By formalising the construction of the QBF proof system obtained from a propositional proof system by adding universal reduction (Beyersdorff, Bonacina & Chew, ITCS '16), we present a new technique for proving proof-size lower bounds in these systems. The technique relies only on two semantic measures: the *cost* of a QBF, and the *capacity* of a proof. By examining the capacity of proofs in several QBF systems, we are able to use the technique to obtain lower bounds based on cost alone. As applications of the technique, we first prove exponential lower bounds for a new family of simple QBFs representing equality. The main application is in proving exponential lower bounds with high probability for a class of randomly generated QBFs, the first 'genuine' lower bounds of this kind, which apply to the QBF analogues of resolution, Cutting Planes, and Polynomial Calculus. Finally, we employ the technique to give a simple proof of hardness for the prominent formulas of Kleine Büning, Karpinski and Flögel.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems: Complexity of proof procedures

Keywords and phrases quantified Boolean formulas, proof complexity, lower bounds

Digital Object Identifier 10.4230/LIPIcs.ITCS.2018.9

1 Introduction

The central question in *proof complexity* can be stated as follows: Given a logical theory and a provable theorem, what is the size of the shortest proof? This question bears tight connections to central problems in computational complexity [19, 25] and bounded arithmetic [42, 24].

Proof complexity is intrinsically linked to recent noteworthy innovations in solving, owing to the fact that any decision procedure implicitly defines a *proof system* for the underlying language. Relating the two fields in this way is illuminating for the practitioner; proof-size and proof-space lower bounds correspond directly to best-case running time and memory consumption for the corresponding solver. Indeed, proof complexity theory has become the main driver for the asymptotic comparison of solving implementations. However, in line with neighbouring fields (such as computational complexity), it is the central task of

[*] This is an extended abstract of the paper available at [7] `https://128.84.21.199/abs/1712.03626`.

9th Innovations in Theoretical Computer Science Conference (ITCS 2018).
Editor: Anna R. Karlin; Article No. 9; pp. 9:1–9:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

demonstrating lower bounds, and of *developing general methods* for showing such results, that proves most challenging for theoreticians.

The desire for general techniques derives from the exceptional strength of modern implementations. Cutting-edge advances in solving, spearheaded by unparalleled progress in Boolean satisfiability (SAT), appear to provide a means for the efficient solution of computationally hard problems [54]. Contemporary SAT solvers routinely dispatch instances in millions of clauses [43], and are effectively employed as **NP**-oracles in more complex settings [44]. The state-of-the-art procedure is based on a propositional proof system called *resolution*, operating on *conjunctive normal form* (CNF) instances using a technique known as *conflict-driven clause learning* (CDCL) [50]. Besides furthering the intense study of resolution and its fragments [19], the evident success has inevitably pushed research frontiers beyond the **NP**-completeness of Boolean satisfiability.

1.1 Beyond propositional satisfiability

A case in point is the logic of *quantified Boolean formulas* (QBF), a theoretically important class that forms the prototypical **PSPACE**-complete language [53]. QBF extends propositional logic with existential and universal quantification, and consequently offers succinct encodings of concrete problems from conformant planning [48, 30, 20], ontological reasoning [40], and formal verification [6], amongst other areas [28, 17, 52]. There is a large body of work on practical QBF solving, and the relative complexities of the associated resolution-type proof systems are well understood [2, 10, 37].

The semantics of QBF has a neat interpretation as a two-player *evaluation game*. Given a QBF $\mathcal{Q} \cdot \phi$, the $\exists$- and $\forall$-players take turns to assign the existential and universal variables of the formula following the order of the quantifier prefix $\mathcal{Q}$. When all variables are assigned, the $\exists$-player wins if the propositional formula ϕ is satisfied; otherwise, the $\forall$-player takes the win. A folklore result states that a QBF is false if and only if the $\forall$-player can win the evaluation game by force; that is, if and only if there exists a winning strategy for the universal player. The concept of *strategy extraction* originates from QBF solving [35], whereby a winning strategy 'extracted' from the proof certifies the truth or falsity of the instance. In practice it is not merely the truth value of the QBF that is required – for real-world applications, certificates provide further useful information [52].

A major paradigm in QBF practice is *quantified conflict-driven clause learning* (QCDCL) [34], a natural extension of CDCL. The vast majority of QBF solvers build upon existing SAT techniques in a similar fashion. Such a notion can hardly be surprising when one considers that an existentially quantified QBF is merely a propositional formula. The novel challenge for the QBF practitioner, therefore, and the real test of a solver's strength, is in the handling of universal quantification.

Proof-theoretic analysis of associated QBF proof systems makes this notion abundantly clear. Consider *QU-Resolution* (QU-Res) [39, 33], a well-studied QBF proof system closely related to QCDCL solving.[1] That calculus simply extends propositional resolution with a *universal reduction* rule, which allows universal literals to be deleted from clauses under certain conditions. On existentially quantified QBFs, therefore, QU-Res is identical to resolution, and proof-size lower bounds for the latter lift immediately to the former. From the viewpoint of quantified logic, lower bounds obtained in this way are rightly considered

[1] The calculus QU-Res, proposed by Van Gelder in [33], generalises Q-Res, introduced by Kleine Büning et al. in [39], by allowing resolution over universally quantified pivots.

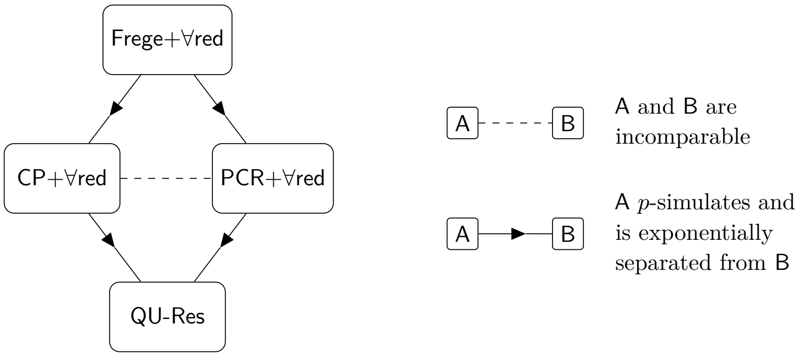

Figure 1 The simulation order of the four QBF proof systems featured in this paper. A proof system A *p*-simulates the system B if each B-proof of a formula Φ can be translated in polynomial time into an A-proof of Φ [25]. If neither A nor B p-simulates the other, then they are incomparable.

non-genuine; they belong in the realm of propositional proof complexity, and tell us nothing about the relative strengths of resolution-based QBF solvers.

Universal reduction is applicable to many suitable propositional proof systems P, giving rise to a general model for QBF systems in the shape of P+∀red [8], which adds to the propositional rules of P the universal reduction rule '∀red'. As a consequence, the phenomenon of genuineness extends well beyond resolution. In this paper, in addition to resolution we consider three stronger systems: Cutting Planes (CP), a well-studied calculus that works with linear inequalities; the algebraic system Polynomial Calculus (with Resolution, PCR); and Frege's eponymous 'textbook' system for propositional logic. Their simulation order is depicted in Figure 1.

What is generally desired (and seemingly elusive) in the QBF community is the development of *general* techniques for *genuine* lower bounds. The current work embraces maximal generality, and contributes a new technique for genuine QBF lower bounds in the general setting of P+∀red.

1.2 When is a lower bound genuine?

Naturally, the aforementioned objections to non-genuine QBF lower bounds may be raised in the abstract setting of P+∀red, as that system encompasses the propositional proof system P. Indeed, given any unsatisfiable propositional formulas that require large proofs in P, one can easily construct any number of contrived QBF families – even with arbitrarily many quantifier alternations – each of which require large proofs in P+∀red, but whose hardness stems from the original propositional formulas. That such lower bounds ought to be identified as non-genuine was highlighted in [21] (cf. also [15]).

The essential point in such cases is that the proofs are large simply because they require many propositional inferences, i.e. many applications of rules of P. Large proofs that do not harbour propositional hardness of this type must therefore contain many universal reductions. Thus, we are brought naturally to a pleasant characterisation of genuine hardness in P+∀red: Genuinely hard QBFs require superpolynomially-many universal reduction steps; all other lower bounds are non-genuine.[2]

[2] This notion can be made formal, as in the oracle model of [15].

In summary, a lower bound on the number of universal reduction steps is always genuine. The technique we introduce in this paper works by counting universal reduction steps, and we therefore deal exclusively in genuine results.

1.3 Random formulas

In the design and testing of solvers, large sets of formulas are needed to make effective comparisons between implementations. While many formulas have been constructed by hand, often representing some combinatorial principle, it is of clear benefit to have a procedure to randomly generate such formulas. The search for a better understanding of when such formulas are likely to be true or false, and their likely hardness for solvers, brings us to the study of the proof complexity of random CNFs and QBFs.

In propositional proof complexity, random 3-SAT instances, the most commonly studied random CNFs, are relatively well understood. There is a constant r such that if a random CNF on n variables contains more than rn clauses, then the CNF is unsatisfiable with probability approaching 1 [32]; the upper bound for r has regularly been improved (see [29], and references therein for previous upper bounds). Further, if the number of clauses is below $n^{6/5-\epsilon}$, the CNF requires exponential-size resolution refutations with high probability [3]. Hardness results for random CNFs are also known for Polynomial Calculus [1, 4] and for Cutting Planes [36, 31].

In contrast, comparatively little is known about randomly generated QBFs. The addition of universally quantified variables raises questions as to what model should be used to generate such QBFs – care is needed to ensure a suitable balance between universal and existential variables.[3] The best-studied model is that of (1,2)-QCNFs [22], for which bounds on the threshold number of clauses needed for a false QBF were shown in [26]. However, to the best of our knowledge, nothing has yet been shown on the proof complexity of randomly generated QBFs. Proving such lower bounds constitutes the major application of our new technique.

1.4 Our contributions

The primary contribution of this work is the proposal of a *novel and semantically-grounded technique* for proving genuine QBF lower bounds in P+∀red, representing a significant forward step in the understanding of reasons for hardness in the proof complexity of quantified Boolean formulas. Our central result, the Size-Cost-Capacity Theorem, provides an *absolute lower bound on the number of universal reductions* for a QBF refutation – in *any* P+∀red proof system – stated as the ratio of two natural measures: the *cost* of a QBF and the *capacity* of a proof. As such, we obtain superpolynomial proof-size lower bounds whenever cost is high and capacity is small.

To that end, we demonstrate that P+∀red proofs have unit capacity when P is resolution or Cutting Planes, and that capacity is at most proof size when P is Polynomial Calculus with Resolution. We therefore obtain lower bounds in these three proof systems based solely on cost. This is a rather pleasant state of affairs, since we are able to apply our technique in three interesting cases simply by demonstating the exponential cost of a family of QBFs. Moerover, in doing so we obtain exponential lower bounds for QU-Res, CP+∀red and PCR+∀red simultaneously.

[3] If any clause contains only universal variables, then there is a constant-size refutation using only this clause.

For our first application, we exemplify our technique with a new family of hard QBFs called the *equality formulas*. We strongly suggest that these formulas, notable for their simplicity and conspicuous exponential cost, will henceforth occupy a prominent place in QBF proof complexity. As our principal application, we prove exponential lower bounds for a large class of randomly generated QBFs by demonstrating that they have high cost with high probability. This is the first time that genuine lower bounds have been shown *en masse* for randomly generated QBFs. As a third example, we show how our technique can be applied to give a simple proof of hardness for the well-known family of QBFs from [39] (cf.[9]).

In addition, we also determine exact conditions on P by which P+$\forall$red is properly defined and receptive to our method, by introducing the notion of a propositional *base system* – a line-based propositional proof system satisfying three natural conditions.

1.5 Organisation of the paper

We continue with the necessary background in Section 2, and provide the details of our P+$\forall$red framework in Section 3. Section 4 presents our lower bound technique, including definitions of cost and capacity, and the statement of our central result, the Size-Cost-Capacity Theorem. Applications of Size-Cost-Capacity, including the details on random formulas, are the subject of Section 5. This is followed in Section 6 by some discussion on the relation of our work to existing QBF techniques, and the merits and future perspectives of our contribution's conceptual innovations. We close the paper in Section 7 with some conclusions and open problems.

2 Preliminaries

2.1 Quantified Boolean formulas

A *conjunctive normal form* (CNF) formula is a conjunction of clauses, each of which is a disjunction of literals. We represent a CNF as a set of clauses, and a clause as a set of literals.

A *quantified Boolean formula* (QBF) in *closed prenex form* is typically denoted $\Phi := \mathcal{Q} \cdot \phi$. In the *quantifier prefix* $\mathcal{Q} := \mathcal{Q}_1 X_1 \cdots \mathcal{Q}_n X_n$, the X_i are pairwise-disjoint sets of Boolean variables (or *blocks*)[4] each of which is quantified either existentially or universally by the *associated quantifier* $\mathcal{Q}_i \in \{\exists, \forall\}$, and consecutive blocks are oppositely quantified. The *propositional part* ϕ is a propositional formula all of whose variables $\mathrm{vars}(\phi)$ are quantified in $\mathcal{Q}$.

By the variables of Φ we mean the set $\mathrm{vars}(\Phi) := \bigcup_{i=1}^{n} X_i$. The set of existential variables of Φ, denoted $\mathrm{vars}_\exists(\Phi)$, is the union of those X_i whose associated quantifier $\mathcal{Q}_i$ is $\exists$, and we define the universal variables of Φ similarly. The prefix $\mathcal{Q}$ imposes a linear order $<_\mathcal{Q}$ on the variables of Φ, such that $x_i <_\mathcal{Q} x_j$ holds whenever $x_i \in X_i$, $x_j \in X_j$ and $i < j$, in which case we say that x_i *is left of* x_j (x_j *is right of* x_i) with respect to $\mathcal{Q}$. We extend the linear order $<_\mathcal{Q}$ to sets of variables in the natural way.

A *literal* l is a Boolean variable x or its negation $\neg x$, and we write $\mathrm{var}(l) := x$. A *total assignment* τ to a set $\mathrm{vars}(\tau) = X$ of Boolean variables is a function $\tau : X \to \{0, 1\}$, typically represented as a set of literals in which the literal $\neg x$ (resp. x) represents the assignment $x \mapsto 0$ (resp. $x \mapsto 1$). The set of all total assignments to X is denoted $\langle X \rangle$. A *partial*

[4] Whereas a block $X = \{x_1, \dots, x_m\}$ is a set, it is written explicitly in a prefix as a string of variables $x_1 \cdots x_m$.

assignment to X is a total assignment to a subset of X. The *projection* of τ to a set X' of Boolean variables is the assignment $\{l \in \tau : \mathrm{var}(l) \in X'\}$.

The *restriction* of Φ by an assignment τ is $\Phi[\tau] := \mathcal{Q}[\tau] \cdot \phi[\tau]$, where $\mathcal{Q}[\tau]$ is obtained from $\mathcal{Q}$ by removing each variable in $\mathrm{vars}(\tau)$ (and any redundant quantifiers), and $\phi[\tau]$ is the restriction of ϕ by τ. Restriction of propositional formulas is defined by the conventional inductive semantics of propositional logic; that is, $\phi[\tau]$ is obtained from ϕ by substituting each occurrence of a variable in $\mathrm{vars}(\tau)$ by its associated truth value, and simplifying the resulting formula in the usual way.

2.2 QBF semantics

Semantics are neatly described in terms of strategies in the two-player *evaluation game*. The game takes place over n rounds, during which the variables of a QBF $\Phi := \mathcal{Q} \cdot \phi$ are assigned strictly in the linear order of the prefix $\mathcal{Q} := \exists E_1 \forall U_1 \cdots \exists E_n \forall U_n$.[5] In the i^{th} round, the existential player selects an assignment α_i to E_i and the universal player responds with an assignment β_i to U_i. At the conclusion the players have constructed a total assignment $\tau := \bigcup_{i=1}^{n}(\alpha_i \cup \beta_i) \in \langle \mathrm{vars}(\Phi) \rangle$. The existential player wins iff $\phi[\tau] = \top$; the universal player wins iff $\phi[\tau] = \bot$.

A strategy for the universal player details exactly how she should respond to all possible moves of the existential player. Formally, a $\forall$-*strategy* for Φ is a function $S : \langle \mathrm{vars}_{\exists}(\Phi) \rangle \to \langle \mathrm{vars}_{\forall}(\Phi) \rangle$ that satisfies the following for each $\alpha, \alpha' \in \mathrm{dom}(S)$ and each $i \in [n]$: if α and α' agree on $E_1 \cup \cdots \cup E_i$, then $S(\alpha)$ and $S(\alpha')$ agree on $U_1 \cup \cdots \cup U_i$.[6] We say that S is *winning* iff $\phi[\alpha \cup S(\alpha)] = \bot$ for each $\alpha \in \mathrm{dom}(S)$.

▶ **Proposition 2.1** (folklore). *A QBF is false if and only if it has a winning $\forall$-strategy.*

2.3 QBF resolution

Resolution is a well-studied refutational proof system for propositional CNF formulas with a single inference rule: the *resolvent* $C_1 \cup C_2$ may be derived from clauses $C_1 \cup \{x\}$ and $C_2 \cup \{\neg x\}$. Resolution is *refutationally* sound and complete: that is, the empty clause can be derived from a CNF iff it is unsatisfiable. Resolution becomes implicationally complete with the addition of the weakening rule, which allows literals to be added to clauses arbitrarily.

QU-Resolution (QU-Res) [39, 33] is a resolution-based proof system for QBFs of the form $\mathcal{Q} \cdot \phi$, where ϕ is a CNF. The calculus supplements resolution with a *universal reduction rule* which allows (literals in) universal variables to be removed from a clause C provided that they are right of all existentials in C with respect to $\mathcal{Q}$. Tautological clauses are explicitly forbidden; for any variable x, one may not derive a clause containing both x and $\neg x$. The rules of QU-Res are given in Figure 2. Note that we choose to include weakening of clauses as a valid inference rule, to emphasize the implicational completeness of the underlying propositional system.

A QU-Res *derivation* of a clause C from Φ is a sequence $C_1, \ldots, C_m$ of clauses in which (a) each C_i is either introduced as an axiom (i.e. $C_i \in \phi$) or is derived from previous clauses in the sequence using resolution or universal reduction, and (b) the *conclusion* $C = C_m$ is the unique clause that is not an antecedent in the application of one of these inference rules. A *refutation* of Φ is a derivation of the empty clause from Φ.

[5] An arbitrary QBF can be written in this form by allowing E_1 and U_n to be empty.
[6] Two assignments agree on a set if and only if their projections to that set are identical.

Axiom:	$\dfrac{}{C}$	C is a clause in the matrix ϕ.
Weakening:	$\dfrac{C}{C \cup W}$	Each variable appearing in W is in $\mathrm{vars}(\varPhi)$. The consequent $C \cup W$ is non-tautologous.
Resolution:	$\dfrac{C_1 \cup \{x\} \qquad C_2 \cup \{\neg x\}}{C_1 \cup C_2}$	The resolvent $C_1 \cup C_2$ is non-tautologous.
Universal reduction:	$\dfrac{C \cup U}{C}$	U contains only universal literals. Each variable in U is right of all existential variables in C, with respect to $\mathcal{Q}$.

Figure 2 The rules of QU-Resolution. The input QBF is $\varPhi = \mathcal{Q} \cdot \phi$, where ϕ is a propositional CNF containing no tautologous clauses.

3 Our framework

3.1 A formal definition of P+∀red

We associate the basic concept of a *line-based propositional proof system* P with the following four features:

(a) A set of *lines* $\mathcal{L}_{\mathsf{P}}$, containing at least the two lines $\top$ and $\bot$ that represent trivial truth and trivial falsity, respectively.

(b) A set of *inference rules* $\mathcal{I}_{\mathsf{P}}$ and an *axiom function* that maps each propositional formula ϕ to a set of axioms $\mathcal{A}_{\mathsf{P}}(\phi) \subseteq \mathcal{L}_{\mathsf{P}}$. The axiom function should be polynomial-time computable, and the validity of inferences should be polynomial-time checkable.

(c) A *variables function* that maps each line $L \in \mathcal{L}_{\mathsf{P}}$ to a finite set of Boolean variables $\mathrm{vars}(L)$, satisfying $\mathrm{vars}(\top) = \mathrm{vars}(\bot) = \emptyset$. Additionally, $\mathrm{vars}(L) \subsetneq \mathrm{vars}(\phi)$ for each line L in a P-derivation from ϕ.[7]

(d) A *restriction operator* (denoted by square brackets) that takes each line $L \in \mathcal{L}_{\mathsf{P}}$, under restriction by any partial assignment τ to $\mathrm{vars}(L)$, to a line $L[\tau] \in \mathcal{L}_{\mathsf{P}}$. If τ is a total assignment, then $L[\tau]$ is either $\top$ or $\bot$.

Universal reduction is a widely used rule of inference in QBF proof systems, by which universal variables may be assigned under certain conditions. More precisely, a line L may be restricted by an assignment to a set of universal variables U provided each $u \in U$ is right of each existential in $\mathrm{vars}(L)$, with respect to the prefix of the input QBF. We state the rule formally in Figure 3.

The primary purpose of universal reduction is to lift a line-based propositional proof system P to QBF, as in the following definition.

▶ **Definition 3.1** (P+∀red [8]). Let P be a line-based propositional proof system. Then P+∀red is the system consisting of the inference rules of P in addition to universal reduction, in which references to the input formula ϕ in the rules of P are interpreted as references to the propositional part of the input QBF $\mathcal{Q} \cdot \phi$.

[7] Note that this does not exclude extended Frege systems (EF), whose lines can be represented as Boolean circuits as in [38, p. 71].

$$\frac{L}{L[\beta]}$$

- β is a partial assignment to the universal variables of Φ.
- each universal in $\mathrm{vars}(\beta)$ is right of each existential in $\mathrm{vars}(L)$, with respect to $\mathcal{Q}$.

Figure 3 The universal reduction rule, where $\Phi = \mathcal{Q} \cdot \phi$ is the input QBF.

The above definition, however, does not gaurantee that P+∀red is sound and complete. To do that, we must work a little harder, and identify some further properties required of P.

Before proceeding, we extend our notation from P to P+∀red in the natural way, denoting the lines available in P+∀red (syntactically equivalent to the lines available in P) by $\mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$, and writing $\mathrm{vars}_\exists(L)$ and $\mathrm{vars}_\forall(L)$ for the subsets of $\mathrm{vars}(L)$ consisting of the variables quantified existentially and universally, with respect to the prefix of the input QBF. Also, we observe that Res + ∀red and QU-Res are (virtually) identical proof systems,[8] and we will henceforth use the latter term.

The size of a P+∀red refutation π, denoted $|\pi|$, is defined similarly as for the propositional system P. (For example, the size of a QU-Res refutation is the number of clauses appearing in it.) For formal definitions of the other propositional systems and their proof sizes, we refer the reader to the full paper.

3.2 Propositional base systems

We first introduce a useful object in our framework: for any line $L \in \mathcal{L}_\mathsf{P}$, an associated Boolean function B_L. Observe that the purpose of the restriction operator is to encompass the natural semantics of P – for that reason, we made the natural stipulation that restriction by a total assignment to the variables of a line yields either trivial truth or trivial falsity. We may therefore associate with any line $L \in \mathcal{L}_\mathsf{P}$ the Boolean function on $\mathrm{vars}(L)$ that computes the propositional models of L, with respect to the semantics of the restriction operator for P.

▶ **Definition 3.2** (associated Boolean function)**.** Let P be a line-based propositional proof system and let $L \in \mathcal{L}_\mathsf{P}$. The *associated Boolean function* for L is $B_L : \langle \mathrm{vars}(L) \rangle \to \{0, 1\}$, defined by $B_L(\tau) = 1$ if $L[\tau] = \top$, and $B_L(\tau) = 0$ otherwise.

Beyond the established notion of 'line-based', we identify three natural conditions on P by which P+∀red is a bona fide, sound and complete QBF proof system. The first of these guarantees that the propositional models of the axioms are exactly those of the input formula, and the second guarantees soundness and completeness *in the classical sense of propositional logic.*[9] The third property ensures that the restriction operator behaves sensibly; that is, the propositional models of the restricted line are computed by the restriction of the associated Boolean function. We introduce the term *base system* for those possessing all three.

▶ **Definition 3.3** (base system)**.** A *base system* P is a line-based propositional proof system satisfying the following three properties:

[8] The only difference between them is that it is allowable to derive universal tautologies and trivial truth in Res + ∀red. Such inferences, however, are never useful.

[9] The (proof-complexity-theoretic) concepts of soundness and completeness for arbitrary proof systems in the sense of Cook and Reckhow are weaker than their counterparts in propositional logic.

- *Axiomatic equivalence.* For each propositional formula ϕ and each $\tau \in \langle \mathrm{vars}(\phi)\rangle$, $\phi[\tau] = \top$ iff each $A \in \mathcal{A}_{\mathsf{P}}(\phi)$ satisfies $A[\tau] = \top$;
- *Inferential equivalence.* For each set of lines $\mathcal{L} \subseteq \mathcal{L}_{\mathsf{P}}$ and each line $L \in \mathcal{L}_{\mathsf{P}}$, L can be derived from $\mathcal{L}$ iff $\mathcal{L}$ semantically entails L;
- *Restrictive closure.* For each $L \in \mathcal{L}_{\mathsf{P}}$ and each partial assignment τ to $\mathrm{vars}(L)$, the Boolean functions $B_{L[\tau]}$ and $B_L|_\tau$ are identical.

On account of the low-level generality, the following theorem requires a non-trivial proof.

▶ **Theorem 3.4.** *If* P *is a base system, then* $\mathsf{P}{+}\forall\mathsf{red}$ *is a sound and complete QBF proof system.*

Formalising the framework of base systems renders our technique applicable to the complete spectrum of $\mathsf{P}{+}\forall\mathsf{red}$ proof systems. All the concrete propositional calculi considered in this work (i.e. those appearing in Figure 1) are demonstrably base systems.

4 Genuine QBF lower bounds with Size-Cost-Capacity

Using an established approach (e.g. [8]), the soundness of $\mathsf{P}{+}\forall\mathsf{red}$ is proved by demonstrating that a winning strategy for the $\forall$-player can be extracted from a refutation. However, with careful construction and analysis of the strategy extraction algorithm, we are able to obtain a much more valuable result – an absolute lower bound on the number of universal reductions steps.

Given a $\mathsf{P}{+}\forall\mathsf{red}$ refutation π of a QBF $\varPhi$, *round-based strategy extraction* works by first restricting π according to the $\exists$-player's move, then collecting the response for the $\forall$-player from some line in π, and iterating until the evaluation game concludes. We therefore reason as follows: A lower bound on the total number of responses contributed by π, coupled with an upper bound on the number of responses contributed per line, yields a lower bound on the number of lines in the refutation. In light of this observation, we define the two measures called *cost* and *capacity*.

4.1 Defining cost

Given a countermodel S for a false QBF $\varPhi$, it is natural to ask how many responses are used for each universal block, since the breadth of responses seems to capture, in some sense at least, the 'size' or 'complexity' of the winning strategy. Let us denote the maximum number (over all universal blocks) of responses in a single block by $\mu(S)$. We contend that $\mu(S)$ is useful measure of a countermodel, with respect to strategy extraction in particular, and so we define the cost of $\varPhi$ as the minimum $\mu(S)$ over all countermodels.

▶ **Definition 4.1** (cost). Let $\varPhi := \forall U_1 \exists E_1 \cdots \forall U_n \exists E_n \cdot \phi$ be a false QBF. Further, for each winning $\forall$-strategy S for $\varPhi$ and each $i \in [n]$, let S_i be the function that maps each $\alpha \in \langle \mathrm{vars}_\exists(\varPhi)\rangle$ to the projection of $S(\alpha)$ to U_i, and let $\mu(S) := \max\{|\mathrm{rng}(S_i)| : i \in [n]\}$. The *cost* of $\varPhi$ is $\mathrm{cost}(\varPhi) := \min\{\mu(S) : S \text{ is a winning } \forall\text{-strategy for } \varPhi\}$.

It should be clear that any winning strategy contains at least $\mathrm{cost}(\varPhi)$ responses to some universal block. With respect to strategy extraction, therefore, cost is a natural semantically-grounded measure that provides a lower bound on the total number of extracted responses.

4.2 Defining capacity

In order to define capacity, we first introduce the concept of a response map. Strictly speaking, given a line $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$ and a total assignment α to the existential variables of L, a response map returns a total assignment to the universal variables that is guaranteed to falsify $L[\alpha]$, as long as such an assignment exists.

▶ **Definition 4.2** (response map). Let P be a base system. A *response map* $\mathcal{R}$ for $\mathsf{P}+\forall\mathsf{red}$ is any function with domain $\{(L, \alpha) : L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}, \alpha \in \langle \text{vars}_\exists(L)\rangle\}$ that maps each (L, α) to some $\beta \in \langle \text{vars}_\forall(L)\rangle$ such that the following holds: If $B_L|_\alpha$ is zero anywhere, then it is zero at β.

Response maps play a vital role in the machinery of strategy extraction in the general setting of $\mathsf{P}+\forall\mathsf{red}$; indeed, for our framework to take effect, it is crucial that strategy extraction can be defined *with respect to an arbitrary response map.*[10] The purpose of capacity, however, is only to provide an upper bound on the number of responses per line. To that end, we define the concept of a *response set* for a line $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$, which is simply a valid set of responses for L according to some response map.

▶ **Definition 4.3** (response set). Let P be a base system, let $\mathcal{R}$ be a response map for $\mathsf{P}+\forall\mathsf{red}$, and let $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$. The set $\{\mathcal{R}(L, \alpha) : \alpha \in \langle \text{vars}_\exists(L)\rangle\}$ is a *response set* for L.

Now, we observe that one may choose to select a response map minimising the size of the response sets for the lines of $\mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$; moreover, round-based strategy extraction returns a winning $\forall$-strategy regardless of the choice of response map. By selecting such a minimal response map $\mathcal{R}$, we will therefore limit the capacity for any line to contribute multiple responses to the extracted strategy. Thus we associate with each P derivation the maximum number of responses that can be extracted from a single line in that derivation, with respect to a minimal response map. This is the intuition behind capacity; it captures the best-case upper bound we can place on the number of responses contributed per line.

▶ **Definition 4.4** (capacity). Let P be a base system, let $\pi = L_1, \ldots, L_m$ be a $\mathsf{P}+\forall\mathsf{red}$ derivation, and let $\mu(L_i) := \min\{|R| : R \text{ is a response set for } L_i\}$, for each $i \in [m]$. The *capacity* of π is $\text{capacity}(\pi) := \max\{\mu(L_i) : i \in [m]\}$.

4.3 The Size-Cost-Capacity Theorem

Putting the two measures together, we obtain our main result, the *Size-Cost-Capacity Theorem.*

▶ **Theorem 4.5** (Size-Cost-Capacity Theorem). *Let P be a base system, and let π be a $\mathsf{P}+\forall\mathsf{red}$ refutation of a QBF Φ. Then*

$$|\pi| \geq \frac{\text{cost}(\Phi)}{\text{capacity}(\pi)}.$$

We emphasize that Size-Cost-Capacity works by counting universal reduction steps, which illustrates that all results obtained by application of our technique are genuine QBF lower bounds in the aforementioned sense.

For the specific applications in this paper, our technique comprises three very useful corollaries of the Size-Cost-Capacity Theorem, obtained in combination with capacity upper

[10] For the details, we kindly refer the reader to the full paper [7].

bounds for specific systems. For example, we prove that all QU-Res and CP+∀red refutations have capacity equal to 1, and hence deduce that *cost alone* gives an absolute proof-size lower bound there.

▶ **Corollary 4.6.** *Let π be a* QU-Res *or* CP+∀red *refutation of a QBF Φ. Then $|\pi| \geq \mathrm{cost}(\Phi)$.*

The case for the QBF version of Polynomial Calculus with Resolution (PCR+∀red) is much more challenging, and requires some linear algebra, owing to the underlying algebraic composition of Polynomial Calculus. Interestingly, it turns out that the capacity of a refutation there is no greater than its size, thus proof size is at least the square root of cost.

▶ **Corollary 4.7.** *Let π be a* PCR+∀red *refutation of a QBF Φ. Then $|\pi| \geq \sqrt{\mathrm{cost}(\Phi)}$.*

Equipped with these results, showing that the cost of a QBF is superpolynomial yields immediate proof-size lower bounds for all three systems simultaneously.

5 Applications of Size-Cost-Capacity

5.1 The equality formulas: a new family of hard QBFs

As a first application of our lower-bound technique, we introduce an interesting new family of hard QBFs.

▶ **Definition 5.1** (equality formulas)**.** For $n \in \mathbb{N}$, the n^{th} *equality formula* is

$$EQ(n) := \exists x_1 \cdots x_n \forall u_1 \cdots u_n \exists t_1 \cdots t_n \cdot \left(\bigwedge_{i=1}^{n} (x_i \vee u_i \vee \neg t_i) \wedge (\neg x_i \vee \neg u_i \vee \neg t_i) \right) \wedge \left(\bigvee_{i=1}^{n} t_i \right).$$

The equality formulas are so called because the only winning strategy for the ∀-player in the evaluation game is as follows: play $u_i = x_i$ for each $i \in [n]$. Consequently the winning strategy is not only unique, it contains all 2^n assignments to the universal variables. These two properties in tandem imply that the equality formulas have exponential cost.

▶ **Proposition 5.2.** *For each $n \in \mathbb{N}$, $\mathrm{cost}(EQ(n)) = 2^n$.*

Applying Size-Cost-Capacity via Corollaries 4.6 and 4.7, we obtain exponential proof-size lower bounds in all three systems QU-Res, CP+∀red and PCR+∀red.

▶ **Theorem 5.3.** *The equality formulas require refutations of size $2^{\Omega(n)}$ in each of the systems* QU-Res*,* CP+∀red *and* PCR+∀red*.*

Whereas it is plausible that the equality formulas are the simplest to which our technique applies, they are without doubt the simplest known hard QBFs. When considering QBF proof complexity lower bounds, particularly in P+∀red systems, we must concern ourselves with formulas with at least a Σ_3 prefix, of which the equality formulas are one of the simplest examples. If a QBF has a Σ_2 prefix, then it is true if and only if the existential parts of the clauses can all be satisfied, i.e. it is equivalent to a SAT problem. Similarly, a refutation of a QBF with a Π_2 prefix consists of a refutation of a subset of the existential clauses corresponding to a particular assignment to the universal variables. A Π_3 formula can also be regarded as essentially a SAT problem using similar reductions as for both Σ_2 and Π_2, so Σ_3 is the smallest prefix where we can expect to find genuine QBF lower bounds.

Closer inspection reveals that this lower bound is of a very specific type – it is a genuine QBF lower bound (the formulas are not harbouring propositional hardness) that does not derive from a circuit lower bound (the winning strategy is not hard to compute in an associated circuit class). In existing QBF literature, the only other example of such a family comes from the famous formulas of Kleine Büning et al. [39] (cf. Subsection 5.3). Those formulas are significantly more complex, and exhibit unbounded quantifier alternation compared to the (bounded) Σ_3 prefix of the equality formulas.

5.2 The first hard random QBFs

For the major application of our technique, we define a class of random QBFs and prove that, with high probability, they are hard in all three systems QU-Res, CP+∀red and PCR+∀red. We generate instances that combine the overall structure of the equality formulas with the literature's existing model of random QBFs [22].

▶ **Definition 5.4.** For each $1 \leq i \leq n$, let $C_i^1, \dots, C_i^{cn}$ be distinct clauses picked uniformly at random from the set of clauses containing 1 literal from the set $X_i := \{x_i^1, \dots, x_i^m\}$ and 2 literals from $Y_i := \{y_i^1, \dots, y_i^n\}$. Define the randomly generated QBF $Q(n,m,c)$ as:

$$Q(n,m,c) := \exists Y_1 \dots Y_n \forall X_1 \dots X_n \exists t_1 \dots t_n \cdot \bigwedge_{i=1}^{n} \bigwedge_{j=1}^{cn} \left(\neg t_i \vee C_i^j\right) \wedge \bigvee_{i=1}^{n} t_i.$$

The specification of how many existential and universal variables each clause should contain is a common and necessary restriction on random QBFs [22, 26]. This prevents the occurrence of a clause containing only universal variables – if such a clause exists, there is a constant size refutation of this clause alone in any P+∀red system. The motivation behind the additional structure in the construction of $Q(n,m,c)$ is that its truth value is equivalent to the disjunction of its 'component parts'; that is $Q(n,m,c) \equiv \bigvee_{i=1}^{n} \Psi_i$, where $\Psi_i := \exists Y_i \forall X_i \cdot \bigwedge_{j=1}^{cn} C_i^j$ for each $i \in [n]$.

These Ψ_i are some of the simplest QBFs one can generate, so $Q(n,m,c)$ is a natural choice for random QBFs. Indeed, the model used to generate the clauses of Ψ_i is also used to generate random formulas for the evaluation of QBF solvers [46, 18].

Drawing on the existing literature [27, 23, 26], we show that suitable choices of the parameters m and c force each Ψ_i to be false with high probability. The individual Ψ_i are essentially equivalent to a random 2-SAT problem, and this step is just an application of results on the satisfiability of such instances.

Moreover, we also prove a cost lower bound. Perhaps surprisingly, this cost lower bound is constructed by applying results on the unsatisfiability of random 2-SAT instances [27] and the truth of random (1,2)-QCNFs [26]. These results both concern only the truth value of the corresponding formulas, and taken individually seem unrelated to cost. However, by carefully choosing the number of clauses so as to allow the application of both results, we can construct a cost lower bound using the following argument.

The Ψ_i are false with high probability, but rearranging the quantifiers to $\forall X_i \exists Y_i \cdot \bigwedge_{j=1}^{cn} C_i^j$ gives a QBF which is true with probability $1 - o(1)$. In other words, with high probability, the universal response in Ψ_i must depend on the existential assignment. That is, it must change depending on the existential assignment, and so with probability $1 - o(1)$, linearly many of the Ψ_i require at least two distinct responses in any winning strategy. By refining our choice of m slightly, this allows us to conclude that $Q(n,m,c)$, with high probability, is a false QBF with large cost.

▶ **Proposition 5.5.** *Let $1 < c < 2$ be a constant, and let $m \leq (1-\epsilon)\log_2(n)$ for some constant $\epsilon > 0$. With probability $1 - o(1)$, $Q(n,m,c)$ is false and* $\mathrm{cost}(Q(n,m,c)) = 2^{\Omega(n^{\epsilon})}$.

Invoking Size-Cost-Capacity yields immediate hardness results. The following theorem constitutes the first proof-size lower bounds for randomly generated formulas in the QBF proof complexity literature. We emphasize that these are genuine QBF lower bounds in the aforementioned sense; they are not merely hard random CNFs lifted to QBF. As for any application of our technique, the refutations are large precisely because they require many universal reduction steps.

▶ **Theorem 5.6.** *Let $1 < c < 2$ be a constant, and let $m \leq (1-\epsilon)\log_2(n)$ for some constant $\epsilon > 0$. With probability $1 - o(1)$, $Q(n, m, c)$ is false, and any* QU-Res*,* CP+∀red *or* PCR+∀red *refutation of $Q(n, m, c)$ requires size $2^{\Omega(n^\epsilon)}$.*

5.3 New proofs of known lower bounds

Our third and final application uses Size-Cost-Capacity to provide a new proof of the hardness of the prominent QBFs of Kleine Büning, Karpinski and Flögel [39]. We consider a common modification of the formulas, denoted by $\lambda(n)$, in which each universal variable is 'doubled'. This modification is known to lift lower bounds from Q-Res to QU-Res [2], where we can apply Size-Cost-Capacity.

By rearranging the quantifier prefix to quantify all the additional universal variables in the penultimate quantifier block, we obtain a cost lower bound for this weaker formula, and so prove the following result.

▶ **Corollary 5.7.** *Any* QU-Res*,* CP+∀red *or* PCR+∀red *proof of $\lambda(n)$ requires size $2^{\Omega(n)}$.*

As QU-Res lower bounds on these modified formulas are shown to be equivalent to Q-Res lower bounds on the original formulas, our technique even proves the original lower bounds from [39] (cf. also [10]), and provides some insight as to the source of hardness.

6 Discussion

6.1 Relation to previous work

It is fair to say that there is a scarcity of general methods for showing genuine lower bounds in systems like P+∀red. In contrast, a number of techniques for propositional calculi have emerged from the intense study of resolution [19, 49].

Researchers have of course attempted to lift these techniques to quantified logic, but with mixed success. The seminal size-width relations for resolution [5], which describe proof size in terms of proof width, are rendered ineffectual by universal quantification [11]. The prover-delayer techniques of [14, 45] have been successfully lifted to QBF, but only apply to the weaker tree-like systems [13], whereas solving techniques such as QCDCL are based on the stronger DAG-like versions. Feasible interpolation [41] is an established propositional technique that has been successfully adapted [12], but it is applicable only to a small class of hand-crafted QBFs of a rather specific syntactic form.

Strategy extraction for QBF lower bounds has been explored previously by exploiting connections to circuit complexity [10, 8, 16]. In particular, [8] established tight relations between circuit and proof complexity, lifting even strong circuit lower bounds for $\mathbf{AC^0}[p]$ circuits [47, 51] to QBF lower bounds for $\mathbf{AC^0}[p]$-Frege+∀red [8], unparalleled in the propositional domain. In fact, for strong proof systems such as Frege+∀red, this strategy extraction technique is sufficient to prove any genuine QBF lower bound, in the sense that any superpolynomial lower bound for Frege+∀red arises either due to a lower bound for Frege, or due to a lower bound for Boolean circuits [16]. However, for weaker systems such as QU-Res, this does not hold; there exist lower bounds which are neither propositional nor circuit lower bounds [15]. The underlying reasons for such hardness results are at present not well understood. A characterisation of such lower bounds, and the proposal of associated lower-bound techniques, would be an important development for QBF proof complexity.

The major drawback of the existing approach of [10, 8, 16], of course, is the rarity of superpolynomial lower bounds from circuit complexity [55], especially for larger circuit classes to which the stronger QBF proof systems connect. With Size-Cost-Capacity we employ a much different approach to strategy extraction. Our technique is motivated by semantics and *does not interface with circuit complexity whatsoever*. Instead, lower bounds are determined directly from the semantic properties of the instance, and consequently we make advances out of the reach of previous techniques.

6.2 Innovations and future perspectives

Our main conceptual innovation is the introduction of *Size-Cost-Capacity*, a semantically-grounded general technique for proving genuine QBF lower bounds.

In this paper, we focus the technique on the P+∀red family of QBF calculi, and prove the first known lower bounds for randomly generated QBFs. The primary appeal of the technique is its semantic nature. We believe that lower bounds based on semantic properties of instances, as opposed to syntactic properties of proofs, work to further our understanding of the hardness phenomenon across the wider range of QBF proof systems. We strongly suggest that Size-Cost-Capacity is applicable beyond P+∀red, and future work will likely establish the hardness of random QBFs in even stronger QBF systems (for example in the expansion based calculus IR-calc [9]).

Size-Cost-Capacity also opens new research avenues concerning the reasons for QBF hardness – a topic that is currently insufficiently understood. Recall that in strong proof systems such as Frege+∀red, superpolynomial proof size lower bounds can be completely characterised: they are either a propositional lower bound or a circuit lower bound [16]. All the QBF families that we consider have no underlying propositional hardness, and winning ∀-strategies can be computed by small circuits, even in very restricted circuit classes. As such, all these QBFs are easy for Frege+∀red.

However, for weaker proof systems, such as QU-Res, CP+∀red and PCR+∀red, propositional hardness and circuit lower bounds alone are not the complete picture. In particular, the lower bounds we show using Size-Cost-Capacity do not fit into either class. That our technique relies on capacity upper bounds which do not hold for strong proof systems leads us to suggest that we have identified a new reason for hardness in those proof systems where the above characterisation does not hold. As such, our work opens the door for a better understanding, and makes steps towards the complete characterisations of reasons for hardness that are currently lacking in the literature.

7 Conclusions

By formalising the conditions on P in the construction of P+∀red, we have developed a new technique for proving QBF lower bounds in P+∀red. The technique depends only on the two natural concepts of the cost of a QBF and the capacity of a proof. Determining the capacity of several well-studied proof systems allowed us to present lower bounds based on cost alone. We have also demonstrated that this technique is not restricted to a few carefully constructed QBFs, but is in fact applicable to a large class of randomly generated formulas, providing the first such lower bound for random QBFs.

References

1 Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. In *Symposium on Foundations of Computer Science (FOCS)*, pages 190–199. IEEE Computer Society, 2001.

2 Valeriy Balabanov, Magdalena Widl, and Jie-Hong R. Jiang. QBF resolution systems and their proof complexities. In Carsten Sinz and Uwe Egly, editors, *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, volume 8561 of *Lecture Notes in Computer Science*, pages 154–169. Springer, 2014.

3 Paul Beame and Toniann Pitassi. Simplified and improved resolution lower bounds. In *Symposium on Foundations of Computer Science (FOCS)*, pages 274–282. IEEE Computer Society, 1996.

4 Eli Ben-Sasson and Russell Impagliazzo. Random CNFs are hard for the polynomial calculus. *Computational Complexity*, 19(4):501–519, 2010.

5 Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.

6 Marco Benedetti and Hratch Mangassarian. QBF-based formal verification: Experience and perspectives. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 5(1-4):133–191, 2008.

7 Olaf Beyersdorff, Joshua Blinkhorn, and Luke Hinde. Size, cost, and capacity: A semantic technique for hard random QBFs. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:35, 2017.

8 Olaf Beyersdorff, Ilario Bonacina, and Leroy Chew. Lower bounds: From circuits to QBF proof systems. In Madhu Sudan, editor, *ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 249–260. ACM, 2016.

9 Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. On unification of QBF resolution-based calculi. In Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik, editors, *International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 8635 of *Lecture Notes in Computer Science*, pages 81–93. Springer, 2014.

10 Olaf Beyersdorff, Leroy Chew, and Mikoláš Janota. Proof complexity of resolution-based QBF calculi. In Ernst W. Mayr and Nicolas Ollinger, editors, *International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 30 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 76–89. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.

11 Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Are short proofs narrow? QBF resolution is not simple. In *Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 15:1–15:14, 2016.

12 Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Feasible interpolation for QBF resolution calculi. *Logical Methods in Computer Science*, 13, 2017.

13 Olaf Beyersdorff, Leroy Chew, and Karteek Sreenivasaiah. A game characterisation of tree-like Q-resolution size. *Journal of Computer and System Sciences*, 2017. in press.

14 Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A characterization of tree-like resolution size. *Information Processing Letters*, 113(18):666–671, 2013.

15 Olaf Beyersdorff, Luke Hinde, and Ján Pich. Reasons for hardness in QBF proof systems. In *Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, 2017.

16 Olaf Beyersdorff and Ján Pich. Understanding Gentzen and Frege systems for QBF. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Symposium on Logic in Computer Science (LICS)*, pages 146–155. ACM, 2016.

17 Roderick Bloem, Robert Könighofer, and Martina Seidl. SAT-based synthesis methods for safety specs. In Kenneth L. McMillan and Xavier Rival, editors, *International Conference*

on Verification, Model Checking, and Abstract Interpretation (VMCAI), volume 8318 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2014.

18 Robert Brummayer, Florian Lonsing, and Armin Biere. Automated testing and debugging of SAT and QBF solvers. In Ofer Strichman and Stefan Szeider, editors, *International Conference on Theory and Practice of Satisfiability Testing (SAT)*, volume 6175 of *Lecture Notes in Computer Science*, pages 44–57. Springer, 2010.

19 Samuel R. Buss. Towards NP-P via proof complexity and search. *Ann. Pure Appl. Logic*, 163(7):906–917, 2012.

20 Michael Cashmore, Maria Fox, and Enrico Giunchiglia. Partially grounded planning as quantified Boolean formula. In Daniel Borrajo, Subbarao Kambhampati, Angelo Oddi, and Simone Fratini, editors, *International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI, 2013.

21 Hubie Chen. Proof complexity modulo the polynomial hierarchy: Understanding alternation as a source of hardness. *ACM Transactions on Computation Theory*, 9(3):15:1–15:20, 2017.

22 Hubie Chen and Yannet Interian. A model for generating random quantified Boolean formulas. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 66–71. Professional Book Center, 2005.

23 Vasek Chvátal and Bruce A. Reed. Mick gets some (the odds are on his side). In *Symposium on Foundations of Computer Science (FOCS)*, pages 620–627. IEEE Computer Society, 1992.

24 Stephen A. Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity*. Cambridge University Press, Cambridge, 2010.

25 Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, 1979.

26 Nadia Creignou, Hervé Daudé, Uwe Egly, and Raphaël Rossignol. Exact location of the phase transition for random (1, 2)-QSAT. *RAIRO - Theoretical Informatics and Applications*, 49(1):23–45, 2015.

27 Wenceslas Fernandez de la Vega. Random 2-SAT: results and problems. *Theoretical Computer Science*, 265(1-2):131–146, 2001.

28 Nachum Dershowitz, Ziyad Hanna, and Jacob Katz. Space-efficient bounded model checking. In Fahiem Bacchus and Toby Walsh, editors, *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, volume 3569 of *Lecture Notes in Computer Science*, pages 502–518. Springer, 2005.

29 Josep Díaz, Lefteris M. Kirousis, Dieter Mitsche, and Xavier Pérez-Giménez. A new upper bound for 3-SAT. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 2 of *LIPIcs*, pages 163–174. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008.

30 Uwe Egly, Martin Kronegger, Florian Lonsing, and Andreas Pfandler. Conformant planning as a case study of incremental QBF solving. *Annals of Mathematics and Artificial Intelligence*, 80(1):21–45, 2017.

31 Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere. Random CNFs are hard for cutting planes. *Computing Research Repository*, abs/1703.02469, 2017.

32 John Franco and Marvin C. Paull. Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem. *Discrete Applied Mathematics*, 5(1):77–87, 1983.

33 Allen Van Gelder. Contributions to the theory of practical quantified Boolean formula solving. In Michela Milano, editor, *International Conference on Principles and Practice of Constraint Programming (CP)*, volume 7514 of *Lecture Notes in Computer Science*, pages 647–663. Springer, 2012.

34 Enrico Giunchiglia, Paolo Marin, and Massimo Narizzano. Reasoning with quantified Boolean formulas. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 761–780. IOS Press, 2009.

35 Alexandra Goultiaeva, Allen Van Gelder, and Fahiem Bacchus. A uniform approach for generating proofs and strategies for both true and false QBF formulas. In Toby Walsh, editor, *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 546–553. IJCAI/AAAI, 2011.

36 Pavel Hrubes and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. *Electronic Colloquium on Computational Complexity*, 24:42, 2017.

37 Mikoláš Janota and João Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theoretical Computer Science*, 577:25–42, 2015.

38 Emil Jeřábek. *Weak pigeonhole principle, and randomized computation.* PhD thesis, Faculty of Mathematics and Physics, Charles University, Prague, 2005.

39 Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. Resolution for quantified Boolean formulas. *Information and Computation*, 117(1):12–18, 1995.

40 Roman Kontchakov, Luca Pulina, Ulrike Sattler, Thomas Schneider, Petra Selmer, Frank Wolter, and Michael Zakharyaschev. Minimal module extraction from DL-lite ontologies using QBF solvers. In Craig Boutilier, editor, *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 836–841. AAAI Press, 2009.

41 Jan Krajíček. Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *Journal of Symbolic Logic*, 62(2):457–486, 1997.

42 Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*, volume 60 of *Encyclopedia of Mathematics and Its Applications.* Cambridge University Press, Cambridge, 1995.

43 Sharad Malik and Lintao Zhang. Boolean satisfiability from theoretical hardness to practical success. *Communications of the ACM*, 52(8):76–82, 2009.

44 Kuldeep S. Meel, Moshe Y. Vardi, Supratik Chakraborty, Daniel J. Fremont, Sanjit A. Seshia, Dror Fried, Alexander Ivrii, and Sharad Malik. Constrained sampling and counting: Universal hashing meets SAT solving. In Adnan Darwiche, editor, *Beyond NP*, volume WS-16-05 of *AAAI Workshops.* AAAI Press, 2016.

45 Pavel Pudlák and Russell Impagliazzo. A lower bound for DLL algorithms for k-SAT (preliminary version). In David B. Shmoys, editor, *Symposium on Discrete Algorithms*, pages 128–136. ACM/SIAM, 2000.

46 Luca Pulina. The ninth QBF solvers evaluation - preliminary report. In Florian Lonsing and Martina Seidl, editors, *International Workshop on Quantified Boolean Formulas (QBF)*, volume 1719 of *CEUR Workshop Proceedings*, pages 1–13. CEUR-WS.org, 2016.

47 Alexander A. Razborov. Lower bounds for the size of circuits of bounded depth with basis $\{\wedge, \oplus\}$. *Mathematical Notes*, 41(4):333–338, 1987.

48 Jussi Rintanen. Asymptotically optimal encodings of conformant planning in QBF. In *National Conference on Artificial Intelligence (AAAI)*, pages 1045–1050. AAAI Press, 2007.

49 Nathan Segerlind. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 13(4):417–481, 2007.

50 João P. Marques Silva and Karem A. Sakallah. GRASP - a new search algorithm for satisfiability. In Rob A. Rutenbar and Ralph H. J. M. Otten, editors, *International Conference on Computer-Aided Design (ICCAD)*, pages 220–227. IEEE Computer Society / ACM, 1996.

51 R. Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In Alfred V. Aho, editor, *ACM Symposium on Theory of Computing (STOC)*, pages 77–82. ACM, 1987.

52 Stefan Staber and Roderick Bloem. Fault localization and correction with QBF. In João Marques-Silva and Karem A. Sakallah, editors, *International Conference on Theory and Applications of Satisfiability Testing (SAT)*, volume 4501 of *Lecture Notes in Computer Science*, pages 355–368. Springer, 2007.

53 Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time: Preliminary report. In Alfred V. Aho, Allan Borodin, Robert L. Constable, Robert W. Floyd, Michael A. Harrison, Richard M. Karp, and H. Raymond Strong, editors, *ACM Symposium on Theory of Computing (STOC)*, pages 1–9. ACM, 1973.

54 Moshe Y. Vardi. Boolean satisfiability: Theory and engineering. *Communications of the ACM*, 57(3):5, 2014.

55 Heribert Vollmer. *Introduction to Circuit Complexity - A Uniform Approach.* Texts in Theoretical Computer Science. Springer, 1999.