

SIZE, COST, AND CAPACITY: A SEMANTIC TECHNIQUE FOR HARD RANDOM QBFs *

OLAF BEYERSDORFF ^a, JOSHUA BLINKHORN ^b, AND LUKE HINDE ^c

^a Friedrich Schiller University Jena, Germany
e-mail address: olaf.beyersdorff@uni-jena.de

^{b,c} School of Computing, University of Leeds, UK
e-mail address: {scjlb,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.

1. INTRODUCTION

1.1. Proof complexity and solving. 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 [Bus12, CR79] and bounded arithmetic [Kra95, CN10].

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

Key words and phrases: quantified Boolean formulas, proof complexity, lower bounds.

ACM subject classification: F.2.2 Nonnumerical Algorithms and Problems: Complexity of proof procedures.

* An extended abstract of this paper appeared in the proceedings of ITCS 2018 [BBH18].

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 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 strength of modern implementations. Cutting-edge advances in solving, spearheaded by progress in Boolean satisfiability (SAT), appear to provide a means for the efficient solution of computationally hard problems [Var14]. Contemporary SAT solvers routinely dispatch instances in millions of clauses [MZ09], and are effectively employed as NP oracles in more complex settings [MVC⁺16]. 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) [SS96]. Besides furthering the intense study of resolution and its fragments [Bus12], the evident success has inevitably pushed research frontiers beyond the NP-completeness of Boolean satisfiability.

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 [SM73]. QBF extends propositional logic with existential and universal quantification, and consequently offers succinct encodings of concrete problems from conformant planning [Rin07, EKLP17, CFG13], ontological reasoning [KPS⁺09], and formal verification [BM08], amongst other areas [DHK05, BKS14, SB07]. 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 [BWJ14, BCJ15, JM15].

The semantics of QBF has a neat interpretation as a two-player *evaluation game*. Given a QBF $Q \cdot \phi$, the \exists - and \forall -players take turns to assign the existential and universal variables of the formula following the order of the quantifier prefix 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 [GGB11], 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 [SB07].

A major paradigm in QBF practice is *quantified conflict-driven clause learning* (QCDCL) [GMN09], 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) [KKF95, Gel12], 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

¹The calculus QU-Res, proposed by Van Gelder in [Gel12], generalises Q-Res, introduced by Kleine Büning et al. in [KKF95], by allowing resolution over universally quantified pivots.

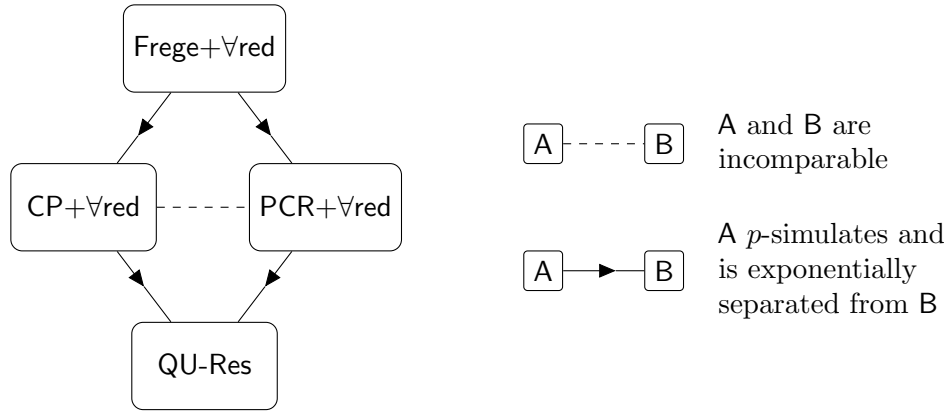


Figure 1: The simulation order of the four QBF proof systems featured in this paper. A proof system A p -simulates the system B if each B -proof of a formula Φ can be translated in polynomial time into an A -proof of Φ [CR79]. If neither A nor B p -simulates the other, then they are incomparable.

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 *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 $P+\forall\text{red}$ [BBC16], which adds to the propositional rules of P the universal reduction rule ‘ $\forall\text{red}$ ’. 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.² The current work embraces maximal generality, and contributes a new technique for genuine QBF lower bounds in the general setting of $P+\forall\text{red}$.

When is a lower bound genuine? Naturally, the aforementioned objections to non-genuine QBF lower bounds may be raised in the abstract setting of $P+\forall\text{red}$, 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 $P+\forall\text{red}$, but whose hardness stems from the original propositional formulas. That such lower bounds ought to be identified as non-genuine was highlighted in [Che17] (cf. also [BHP17]).

²Instead of using a more descriptive name such as ‘quantifier-alternation-exploiting lower bounds’ we refer throughout this work simply to ‘genuine lower bounds’ to preserve readability, hopefully without loss in clarity.

A formal model for genuine QBF lower bounds was proposed in [BHP17]. The model essentially adds an NP oracle to the proof system, which allows any logically correct propositional inference to be derived in a single step. Propositional hardness is thus removed from the system. In the context of $P+\forall\text{red}$, the addition of an NP oracle essentially gives P-derivations for free; any formula derivable in P can be introduced immediately. If a lower bound persists despite the oracle, then the proof either requires many universal reduction steps, or requires reductions on large lines.

The model provides an adequate definition of a genuine lower bound, and we adopt a similar approach in this article. We work with *semantic* $P+\forall\text{red}$ refutations, in which any logically correct P inference or universal reduction step is allowed (Definition 4.5). The technique we introduce gives a lower bound on the size of a semantic $P+\forall\text{red}$ refutation, i.e. a lower bound that still holds with an NP oracle. We therefore deal exclusively in genuine results.

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 [FP83]; the upper bound for r has regularly been improved (see [DKMP08], 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 [BP96]. Hardness results for random CNFs are also known for Polynomial Calculus [AR01, BI10] and for Cutting Planes [HP17, FPPR17].

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 [CI05], for which bounds on the threshold number of clauses needed for a false QBF were shown in [CDER15]. 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.

1.2. 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 $P+\forall\text{red}$, representing a significant forward step in the understanding of reasons for hardness in the proof complexity of quantified Boolean formulas.

We exemplify the technique with a new family of hard QBFs, notable for their simplicity, which we strongly suggest should henceforth occupy a prominent place in QBF proof complexity. As our principal application we prove exponential lower-bounds in three

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

concrete $P+\forall\text{red}$ systems for a large class of randomly generated QBFs. This is the first time that genuine lower bounds have been shown *en masse* for randomly generated QBFs. Lastly, we note that our technique can be applied to give a simple proof of hardness for a family of well known QBFs from [KKF95].

In addition, we also determine exact conditions on a so-called *base system* P by which $P+\forall\text{red}$ is properly defined and receptive to our method. We detail our contributions below, beginning with a description of the framework, followed by the lower bound technique, and concluding with several applications.

1.2.1. *The universal reduction formalism.* In order to present our contribution in total generality, we take the concept of $P+\forall\text{red}$ (introduced in [BBC16] for a hierarchy of Frege systems) and formalise conditions on P yielding a sound and complete QBF proof system (Theorem 4.7). We identify three natural properties that are sufficient: (a) The set of derivable axioms is semantically equivalent to the input formula; (b) The system exhibits logical correctness and implicational completeness; (c) The system is closed under restrictions. Any line-based propositional calculus possessing all three properties is referred to as a *base system* (Definition 4.2). Formalising the framework of base systems renders our technique applicable to the complete spectrum of $P+\forall\text{red}$ systems. All the concrete propositional calculi considered in this work (i.e. those appearing in Figure 1) are demonstrably base systems. We note that static propositional proof systems (such as Nullstellensatz) cannot be upgraded to $P+\forall\text{red}$ simply because they are not line-based.

1.2.2. *A new technique for genuine QBF lower bounds.* Our technique is based on careful construction and analysis of *strategy extraction* in $P+\forall\text{red}$ [BBC16], a process by which a winning strategy for the universal player can be efficiently obtained from a refutation. Given a $P+\forall\text{red}$ refutation π of a QBF Φ , *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 define two measures called *cost* and *capacity* (Definitions 5.1 and 5.6). The cost of Φ is defined such that any winning strategy contains at least $\text{cost}(\Phi)$ responses to some universal block. Cost, therefore, is a natural semantically-grounded measure that provides a lower bound on the total number of extracted responses. The upper bound is given by the capacity of π , a measure defined such that any response contributed from a given line in π may be selected from a set of cardinality at most $\text{capacity}(\pi)$.

Putting the two measures together, we obtain our main result, the *Size-Cost-Capacity Theorem* (Theorem 5.12), which states that the size of π is at least $\text{cost}(\Phi)/\text{capacity}(\pi)$. We also show explicitly that Size-Cost-Capacity returns a lower bound on the size of a *semantic* $P+\forall\text{red}$ refutation, which illustrates that all results obtained by application of our technique are genuine QBF lower bounds in the aforementioned sense.

For direct applications of Size-Cost-Capacity, we first supply upper bounds on the capacity of refutations in concrete $P+\forall\text{red}$ proof systems. We prove that all QU-Res and CP $+\forall\text{red}$ refutations have capacity equal to 1 (Propositions 5.7 and 6.2), whereupon *cost alone* gives an absolute proof-size lower bound (Corollaries 5.13 and 6.3). The case for the QBF version of Polynomial Calculus with Resolution (PCR $+\forall\text{red}$) 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 in that

system is no greater than its size (Proposition 6.5), thus proof size is at least the square root of cost (Corollary 6.6). Equipped with these corollaries, showing that the cost of a QBF is superpolynomial yields immediate proof-size lower bounds for all three systems simultaneously.

1.2.3. *Applications of the technique.* We demonstrate the effectiveness of our new technique on three applications.

A. The equality formulas: a non-trivial special case. We introduce a new family of hard QBFs that we call the *equality formulas* (Definition 3.1), so called because the only winning strategy for the \forall -player in the evaluation game is to copy exactly the moves of the \exists -player. We first prove that the equality formulas have exponential cost (Proposition 5.2). Using Size-Cost-Capacity, we therefore prove that they require exponential-size refutations in QU-Res, CP+ \forall red and PCR+ \forall red. We also demonstrate that the formulas have linear-size refutations in Frege+ \forall red (Proposition 6.7), which shows that Frege+ \forall red proofs can have exponential capacity.

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. [KKF95] (cf. item C. below). Those formulas are significantly more complex, and exhibit unbounded quantifier alternation compared to the (bounded) Σ_3 prefix of the equality formulas.

Indeed, the equality formulas appear to capture rather well the role that high cost plays in round-based strategy extraction to enforce large refutations. For that reason, we include a direct proof of hardness for QU-Res in Section 3. This intended as a primer, to illustrate by example the concept and method-of-proof behind our lower bound technique.

B. The first hard random QBFs. For the major application of our technique, we define a class of random QBFs (Definition 7.1) and prove that they are hard with high probability in all three systems QU-Res, CP+ \forall red and PCR+ \forall red (Theorem 7.9). We generate instances that combine the overall structure of the equality formulas with the existing model of random QBFs [CI05] used in the competitive evaluation of solvers [Pul16, BLB10]. Thus, while at first glance these formulas may seem rather structured, they are simply a disjunction of QBFs generated using a minor variation on the (1,2)-QCNF model of [CI05].

Drawing on the existing literature [div01, CR92, CDER15], we show that suitable choices of parameters force our generated formulas to be false and have exponential cost with high probability (Lemma 7.8). Perhaps surprisingly, the cost lower bound is constructed by applying results on the unsatisfiability of random 2-SAT instances [div01] and the truth of random (1,2)-QCNFs [CDER15]. 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 a novel semantic argument.

Our contribution 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.

C. New proofs of known lower bounds. Our final application uses Size-Cost-Capacity to provide a new proof of the hardness of the prominent QBFs of Kleine Büning, Karpinski and Flögel [KKF95]. 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 in Q-Res to lower bounds in QU-Res [BWJ14], 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 weaker formula with exponential cost (Proposition 8.2). An elementary reduction then shows that $\lambda(n)$ requires exponential size QU-Res refutations (Corollary 8.3). 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 [KKF95] (cf. also [BCJ15]), and provides some insight as to the source of hardness.

1.3. 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 $P+\forall\text{red}$. In contrast, a number of techniques for propositional calculi have emerged from the intense study of resolution [Bus12, Seg07].

Researchers have of course attempted to lift these techniques to quantified logic, but with mixed success. The seminal size-width relations for resolution [BW01], which describe proof size in terms of proof width, are rendered ineffectual by universal quantification [BCMS18a]. The prover-delayer techniques of [BGL13, PI00] have been successfully lifted to QBF, but only apply to the weaker tree-like systems [BCS17], whereas solving techniques such as QCDCL are based on the stronger DAG-like versions. Feasible interpolation [Kra97] is an established propositional technique that has been successfully adapted [BCMS17], 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 [BCJ15, BBC16, BP16]. In particular, [BBC16] established tight relations between circuit and proof complexity, lifting even strong circuit lower bounds for $AC^0[p]$ circuits [Raz87, Smo87] to QBF lower bounds for $AC^0[p]-\text{Frege}+\forall\text{red}$ [BBC16], which is unparalleled in the propositional domain. In fact, for strong proof systems such as $\text{Frege}+\forall\text{red}$, this strategy extraction technique is sufficient to prove any genuine QBF lower bound, in the sense that any superpolynomial lower bound for $\text{Frege}+\forall\text{red}$ arises either due to a lower bound for Frege , or due to a lower bound for Boolean circuits [BP16]. However for weaker systems such as QU-Res, this does not hold and there exist lower bounds which are neither a propositional lower bound nor a circuit lower bound [BHP17]. The reasons underlying such hardness results are at present not well understood. The development of techniques for, or a characterisation of, such lower bounds would be an important step in QBF proof complexity.

The major drawback of the existing approach of [BCJ15, BBC16, BP16], of course, is the rarity of superpolynomial lower bounds from circuit complexity [Vol99], 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.

1.4. Organisation of the paper. We provide the relevant background in Section 2. In Section 3, we introduce the equality formulas and give the direct proof of hardness for QU-Res, while providing an overview of round-based strategy extraction. Section 4 introduces our framework, including the formal definition of a base system and the proofs of soundness and completeness of P+ \forall red. Our lower bound technique follows in Section 5, comprising definitions of cost and capacity and the proof of the Size-Cost-Capacity Theorem. Upper bounds on capacity for CP+ \forall red and PCR+ \forall red are the subject of Section 6. Our major application to random QBFs is presented in Section 7. In Section 8 we apply Size-Cost-Capacity to the formulas of Kleine Büning et al. [KKF95]. We offer some concluding thoughts in Section 9.

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, and a *literal* is a Boolean variable or its negation. 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 $\text{vars}(\phi)$ are quantified in \mathcal{Q} . A *QCNF* is a QBF whose propositional part is a CNF.

For a literal l , we write $\text{var}(l) := x$ iff $l = x$ or $l = \neg x$. By the variables of Φ we mean the set $\text{vars}(\Phi) := \bigcup_{i=1}^n X_i$. The set of existential variables of Φ , denoted $\text{vars}_{\exists}(\Phi)$, 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} defines a binary relation $<_{\mathcal{Q}}$ on the variables of Φ , such that $x_i <_{\mathcal{Q}} x_j$ holds iff $x_i \in X_i$, $x_j \in X_j$ 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} . For two variable sets $X, X' \subseteq \text{vars}(\Phi)$, we write $X <_{\mathcal{Q}} X'$ iff each variable in X is left of each variable in X' with respect to \mathcal{Q} .

A *literal* l is a Boolean variable x or its negation $\neg x$, and we write $\text{var}(l) := x$. A *total assignment* τ to a set $\text{vars}(\tau) = X$ of Boolean variables is a function $\tau : X \rightarrow \{0, 1\}$, 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 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 $\{l \in \tau : \text{var}(l) \in X'\}$.

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 $\text{vars}(\tau)$ (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

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

each occurrence of a variable in $\text{vars}(\tau)$ 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 $\mathcal{Q} := \exists E_1 \forall U_1 \dots \exists E_n \forall U_n$.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 $\tau := \bigcup_{i=1}^n (\alpha_i \cup \beta_i) \in \langle \text{vars}(\Phi) \rangle$. The existential player wins iff $\phi[\tau] = \top$; the universal player wins iff $\phi[\tau] = \perp$.

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 $S : \langle \text{vars}_{\exists}(\Phi) \rangle \rightarrow \langle \text{vars}_{\forall}(\Phi) \rangle$ that satisfies the following for each $\alpha, \alpha' \in \text{dom}(S)$ and each $i \in [n]$: if α and α' agree on $E_1 \cup \dots \cup E_i$, then $S(\alpha)$ and $S(\alpha')$ agree on $U_1 \cup \dots \cup U_i$.⁶ We say that S is *winning* iff $\phi[\alpha \cup S(\alpha)] = \perp$ for each $\alpha \in \text{vars}_{\exists}(\Phi)$.

Existential strategies are defined dually. Given a QBF $\Phi' := \forall U_1 \exists E_1 \dots \forall U_n \exists E_n \cdot \phi$, an \exists -*strategy* for Φ' is a function $S' : \langle \text{vars}_{\forall}(\Phi') \rangle \rightarrow \langle \text{vars}_{\exists}(\Phi') \rangle$ that satisfies the following for each $\beta, \beta' \in \text{dom}(S')$ and each $i \in [n]$: if β and β' agree on $U_1 \cup \dots \cup U_i$, then $S(\beta)$ and $S(\beta')$ agree on $E_1 \cup \dots \cup E_i$. We say that S is *winning* iff $\phi[\beta \cup S'(\beta)] = \top$ for each $\alpha \in \text{vars}_{\forall}(\Phi')$.

Proposition 2.1 (folklore). *A QBF is false if and only if it has a winning \forall -strategy, and is true if and only if it has winning \exists -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 implicational complete with the addition of the weakening rule, which allows literals to be added to clauses arbitrarily.

QU-Resolution (QU-Res) [KKF95, Gel12] is a resolution-based proof system for false QCNFs. 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 a QCNF Φ is a sequence C_1, \dots, 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 .

⁵An arbitrary QBF can be written in this form by allowing E_1 and U_n to be empty.

⁶Two assignments agree on a set if and only if their projections to that set are identical.

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

Figure 2: The rules of QU-Resolution. The input is a QCNF $\Phi = \mathcal{Q} \cdot \phi$ whose propositional part contains no tautological clauses.

3. A DIRECT PROOF OF HARDNESS FOR THE EQUALITY FORMULAS

In this section, we introduce the *equality formulas* and give a direct proof of their hardness in the well-known QBF proof system QU-Res. The material in this section is intended to illuminate, by means of an accessible example, the paradigm of round-based strategy extraction, and our exploitation of it as a new lower-bound technique.

3.1. The equality formulas. The salient feature of the equality formulas, defined below, is that each instance has a unique winning strategy, and the cardinality of its range is exactly 2^n .

Definition 3.1 (equality formulas). For $n \in \mathbb{N}$, the n^{th} equality formula is

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

Note that the propositional part of $EQ(n)$ is the CNF consisting of the *long clause* $\{t_1, \dots, t_n\}$ and each pair of clauses $\{x_i, u_i, \neg t_i\}, \{\neg x_i, \neg u_i, \neg t_i\}$ for $i \in [n]$.

The equality formulas are false, and it is clear that there is only one winning strategy for the universal player; namely, she must assign each u_i the same value as the corresponding x_i . Proceeding this way, she forces all n unit clauses $\{\neg t_i\}$ to be present on the board with only the final block left to play. Then the existential player must lose, since satisfying all such unit clauses entails falsifying the long clause $\{t_1, \dots, t_n\}$. This is indeed the only way to win, since any other reply from the universal player would drop at least one unit clause, allowing her opponent to satisfy the long clause.

The upshot is that the existential player can *force* his opponent to play any one of the total assignments to the universal variables. It follows that the range of the unique winning \forall -strategy for $EQ(n)$ is exactly the set $R := \langle \{u_1, \dots, u_n\} \rangle$. The fact that this set has exponential cardinality is a key feature that we exploit in our lower bound proof.

3.2. Overview of round-based strategy extraction. Strategy extraction is an important QBF paradigm that was motivated by solving certification (cf. [NPPT09, GGB11]), and subsequently received much attention in the literature [BJ12, PSS16, ELW13, BBC16]. In this paper, we follow the algorithm given in [GGB11], which for the sake of clarity we refer to as *round-based strategy extraction*.

Given a QU-Res refutation of a QCNF $\exists E_1 \forall U_1 \dots \exists E_n \forall U_n \cdot \phi$, round-based strategy extraction is an iterative procedure that computes a winning \forall -strategy. During the course of the game, the \forall -player maintains a restriction of the refutation, from which her winning responses may be determined. In each round she performs two operations:

- restrict the current refutation by the \exists -player's move $\alpha_i \in \langle E_i \rangle$;
- ‘read off’ a response $\beta_i \in \langle U_i \rangle$ from the current refutation, then restrict by β .

These operations are repeated round by round until the game concludes. At termination, we obtain a refutation of the input formula under a total assignment. Since QU-Res is sound, the assignment must falsify the propositional part of the formula. Hence, the correctness of the procedure rests on the fact that restrictions by α_i and β_i preserve the refutation. This is detailed below over Propositions 3.2 and 3.3.

To streamline the material, we work with a normal form of QU-Res derivation in which universal reduction has stronger side conditions: one must remove all the universal literals from the rightmost universal block appearing in the clause. Formally, one may derive C from $C \cup R$ by universal reduction provided that $\text{vars}(C) <_{\mathcal{Q}} \text{vars}(R)$ and $\text{vars}(R)$ is a subset of some universal block U in \mathcal{Q} , where \mathcal{Q} is the prefix of the input QBF.

It is easy to see that QU-Res derivations can be placed in such a normal form with no increase in size. To do so, one first applies as large a universal reduction as possible to every clause, yielding a new refutation in which each clause is a subset of the corresponding one from the original refutation. One then takes the first occurrence of the empty clause as the conclusion, and keeps only the subderivation of this conclusion, discarding all other clauses (as well as any duplicates), which are rendered redundant.

For the rest of this section, we assume QU-Res derivations to be in this normal form. It is easy to see that normal refutations need *at most one universal* reduction on the leftmost block, since the consequent of such a reduction must be the empty clause.

Restriction by the \exists -player's move. Informally, the restriction of a QU-Res refutation π by an arbitrary assignment δ (denoted $\pi[\delta]$) is obtained simply by restricting the axiom clauses by δ , discarding any that are satisfied, and carrying out the steps of the original refutation wherever possible. Under a formal definition of restriction, it is easy to prove that existential restrictions preserve QU-Res refutations.

Proposition 3.2. *Let π be a QU-Res refutation of a QCNF Φ , and let α be a partial assignment to the existential variables of Φ . Then $\pi[\alpha]$ is a QU-Res refutation of $\Phi[\alpha]$.*

It follows immediately that restriction by the \exists -player's move $\alpha_i \in \langle E_i \rangle$ preserves the current refutation.

Restriction by the \forall -player's response. Arbitrary universal restrictions do not preserve QU-Res refutations, as it is possible that a universal assignment satisfies the antecedent of a universal reduction step, but not the consequent. However, a universal restriction does indeed preserve a refutation provided it does not satisfy any reduced literals. If the first

block U of \mathcal{Q} is universal, then we have at most one universal reduction on U , since the refutation finishes after this step. Therefore, restriction by a total assignment to U need only be consistent with the literals reduced in this step (if it exists) in order to preserve the refutation. For non-trivial refutations, all of these reduced literals appear in the penultimate clause.

Proposition 3.3. *Let π be a non-trivial QU-Res refutation of a QCNF Φ whose first block U is universal, let $\beta \in \langle U \rangle$, and let C be the set of U -literals appearing in the penultimate clause of π . If $\{\neg l : l \in C\} \subseteq \beta$, then $\pi[\beta]$ is a refutation of $\Phi[\beta]$.*

Hence, if the current refutation is non-trivial, the \forall -player may read off the U -literals from the penultimate clause, negate them, and extend them arbitrarily to an assignment $\beta_i \in \langle U_i \rangle$. Preservation of the current refutation under restriction by β_i is guaranteed by Proposition 3.3. If the current refutation is trivial, an arbitrary β_i is clearly sufficient.

3.3. Direct proof of hardness. Propositions 3.2 and 3.3 form the basis of a lower-bound proof for the equality formulas, via the following lemma.

Lemma 3.4. *Let $n \in \mathbb{N}$, let π be a QU-Res refutation of $EQ(n)$, and let β be a total assignment to the universal variables of $EQ(n)$. Then there exists a clause in π that contains $\{\neg l : l \in \beta\}$ as a subset.*

Proof. Let $X := \{x_1, \dots, x_n\}$, let $U := \{u_1, \dots, u_n\}$, and let α be the total assignment to X defined by

$$\alpha(x_i) := \beta(u_i), \text{ for each } i \in [n].$$

Observe that, by the nature of the equality formulas, β is the unique total assignment to U under which $EQ(n)[\alpha]$ is false. Moreover, $\pi[\alpha]$ is non-trivial, since the propositional part of $EQ(n)[\alpha]$ does not contain the empty clause. Let C be the set of U -literals appearing in the penultimate clause of $\pi[\alpha]$, and let $\beta' \in \langle U \rangle$. Now, if $\{\neg l : l \in C\} \subseteq \beta'$, then $EQ(n)[\alpha \cup \beta']$ is false, by Proposition 3.3 and the soundness of QU-Res. Therefore, if β' extends $\{\neg l : l \in C\}$, then $\beta' = \beta$. It follows that $\beta = \{\neg l : l \in C\}$, or equivalently, $C = \{\neg l : l \in \beta\}$. The lemma follows since any clause in $\pi[\alpha]$ is a subset of some clause in π , by definition of restriction of clauses. \square

Since QU-Res disallows tautological clauses, the lower bound is an easy consequence of the preceding lemma.

Theorem 3.5. *Let $n \in \mathbb{N}$, and let π be a QU-Res refutation of $EQ(n)$. Then $|\pi| \geq 2^n$.*

Proof. Aiming for contradiction, suppose that $|\pi| < 2^n$. Observe that there are exactly 2^n total assignments to $\{u_1, \dots, u_n\}$. Then, by Lemma 3.4, there exist distinct total assignments $\beta_1, \beta_2 \in \langle \{u_1, \dots, u_n\} \rangle$ and a clause C in π such that $\{\neg l : l \in \beta_1\}$ and $\{\neg l : l \in \beta_2\}$ are both subsets of C . Since β_1 and β_2 are distinct, there is some literal l satisfying $l \in \beta_1$ and $\neg l \in \beta_2$. But then C is a tautological clause, and does not appear in any QU-Res derivation. \square

In a nutshell, Theorem 3.5 was proved by equating the minimum refutation size with the cardinality of the range of a winning \forall -strategy for $EQ(n)$. Our argument here was aided by two facts: $EQ(n)$ has a unique winning \forall -strategy and contains a single universal block. Of course, neither fact holds for QBFs in general. Nonetheless, in the sequel we generalise the method to prove an absolute proof-size lower bound for any instance in $P+\forall$ red.

4. OUR FRAMEWORK

In this section, we develop a framework for Size-Cost-Capacity centred on a precise description of the $P+\forall\text{red}$ formalism. In Subsection 4.1, we first describe what we mean by a ‘line-based’ propositional proof system P , and then introduce the notion of a ‘base system’ - a line-based system satisfying some natural conditions. In Subsection 4.2 we define $P+\forall\text{red}$ with respect to a line-based system P , and prove that it is sound and complete if P is a base system.

4.1. Propositional base systems. We associate the basic concept of a *line-based propositional proof system* P with the following two features:

- A set of *lines* \mathcal{L}_P , containing at least the two lines \top and \perp that represent trivial truth and trivial falsity, respectively.
- A set of *inference rules* \mathcal{I}_P and an *axiom function* that maps each propositional formula ϕ to a set of axioms $\mathcal{A}_P(\phi) \subseteq \mathcal{L}_P$. The axiom function should be polynomial-time computable, and the validity of applications of inference rules should be polynomial-time checkable.

Following convention, a P -derivation from a propositional formula ϕ is a sequence $\pi = L_1, \dots, L_m$ of lines from \mathcal{L}_P , in which each line L_i is either an axiom from the set $\mathcal{A}_P(\phi)$, or may be derived from previous lines using an inference rule in \mathcal{I}_P . The final line L_m is called the *conclusion* of π , and π is a refutation iff $L_m = \perp$.

In order to facilitate the restriction of P -derivations, we require two further features:

- A *variables function* that maps each line $L \in \mathcal{L}_P$ to a finite set of Boolean variables $\text{vars}(L)$, satisfying $\text{vars}(\top) = \text{vars}(\perp) = \emptyset$. Additionally, $\text{vars}(L) \subseteq \text{vars}(\phi)$ for each line L in a P -derivation from ϕ .⁷
- A *restriction operator* (denoted by square brackets) that takes each line $L \in \mathcal{L}_P$, under restriction by any partial assignment τ to $\text{vars}(L)$, to a line $L[\tau] \in \mathcal{L}_P$. If τ is a total assignment, then $L[\tau]$ is either \top or \perp . Restriction of L by an arbitrary Boolean assignment σ is defined as the restriction of L by the projection of σ to $\text{vars}(L)$.

It should be clear 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 $L \in \mathcal{L}_P$ the Boolean function on $\text{vars}(L)$ that computes the propositional models of L , with respect to the semantics of the restriction operator for P .

Definition 4.1 (associated Boolean function). Let P be a line-based propositional proof system and let $L \in \mathcal{L}_P$. The *associated Boolean function* for L is $B_L : \langle \text{vars}(L) \rangle \rightarrow \{0, 1\}$, defined by

$$B_L(\tau) = \begin{cases} 1, & \text{if } L[\tau] = \top, \\ 0, & \text{if } L[\tau] = \perp. \end{cases}$$

It is useful to define the usual notion of semantic entailment on the lines of P . Given two lines $L, L' \in \mathcal{L}_P$, we say that L' *semantically entails* L if

$$L'[\tau] = \top \quad \Rightarrow \quad L[\tau] = \top, \quad \text{for each } \tau \in \langle \text{vars}(L) \cup \text{vars}(L') \rangle,$$

⁷Note that this does not exclude extended Frege systems (EF), whose lines can be represented as Boolean circuits as in [Jer05, p. 71].

and L and L' are said to be *semantically equivalent* (written $L \equiv L'$) if they semantically entail one another. We say that a finite set $\{L_1, \dots, L_k\} \subseteq \mathcal{L}_{\mathcal{P}}$ of lines semantically entails L if

$$L_i[\tau] = \top \text{ for each } i \in [k] \quad \Rightarrow \quad L[\tau] = \top, \quad \text{for each } \tau \in \left\langle \bigcup_{i \in [n]} \text{vars}(L_i) \cup \text{vars}(L) \right\rangle.$$

Beyond the notion of *line-based*, we identify three natural properties. 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*.⁸ 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 4.2 (base system). A *base system* \mathcal{P} is a line-based propositional proof system satisfying the following three properties:

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

As a first example, we note that resolution (with weakening, and excluding tautological clauses) forms a base system. The axiomatic equivalence is trivial, as is inferential equivalence, which follows directly from implicational completeness⁹ and logical correctness. Taking the conventional definitions of the variable function and restriction operator, the restrictive closure of resolution is readily verified. This is to be expected of course, since the restriction of clauses is based on a standard definition of semantics in propositional logic.

4.2. The $\mathcal{P}+\forall\text{red}$ formalism. 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 a partial assignment to a universal block U provided it is *rightmost* in $\text{vars}(L)$, with respect to the prefix \mathcal{Q} of the input QBF; that is, U is right of every block in \mathcal{Q} whose intersection with $\text{vars}(L)$ is non-empty. We state the rule formally in Figure 3. By restrictive closure, the restriction of a line by an assignment to β results in the exclusion of $\text{vars}(\beta)$ from the domain of the associated Boolean function. Universal reduction should therefore be viewed as a sound method for deleting universal variables.

The purpose of universal reduction is to lift a propositional proof system \mathcal{P} to a QBF system $\mathcal{P}+\forall\text{red}$, as in the following definition.

Definition 4.3 ($\mathcal{P}+\forall\text{red}$ [BBC16]). Let \mathcal{P} be a line-based propositional proof system. Then $\mathcal{P}+\forall\text{red}$ is the system consisting of the inference rules of \mathcal{P} in addition to universal reduction, in which references to the input formula ϕ in the rules of \mathcal{P} are interpreted as references to the propositional part of the input QBF $\mathcal{Q} \cdot \phi$.

⁸The (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.

⁹Formally, one should introduce a special clause \top that can be derived from the empty set of clauses.




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

We lift some notation from P to $\mathsf{P}+\forall\text{red}$ as follows. We denote the lines available in $\mathsf{P}+\forall\text{red}$ (syntactically equivalent to the lines available in P) by $\mathcal{L}_{\mathsf{P}}^{\mathcal{Q}}$, where \mathcal{Q} is the prefix of the input QBF. For $L \in \mathcal{L}_{\mathsf{P}}^{\mathcal{Q}}$, we write $\text{vars}_{\exists}(L)$ and $\text{vars}_{\forall}(L)$ for the subsets of $\text{vars}(L)$ consisting of the existentially and universally quantified variables, respectively.

Soundness and completeness of $\mathsf{P}+\forall\text{red}$ is not guaranteed for an arbitrary line-based system P , but it is guaranteed if P is a base system. The following lemma establishes completeness.

Lemma 4.4. *Let P be a base system. Every false QBF has a $\mathsf{P}+\forall\text{red}$ refutation.*

Proof. Let $\Phi = \mathcal{Q} \cdot \phi$ be a false QBF with universal variables $\text{vars}_{\forall}(\Phi) = \{u_1, \dots, u_n\}$, and let S be a winning \forall -strategy for Φ . For each $i \in [n]$, let X_i be the set of existential variables left of u_i in \mathcal{Q} .

For each $i \in [n]$, let B_i be the Boolean function with domain $\langle X_i \cup \{u_1, \dots, u_i\} \rangle$ that maps to 0 iff, for each $j \in [i]$, the assignment of u_j matches that of the strategy on the assignment to X_j . Formally, $B_i(\sigma) = 0$ iff $\sigma_{\forall} \subseteq S(\sigma'_{\exists})$ for every extension of σ_{\exists} to $\text{vars}_{\exists}(\Phi)$, where σ_{\exists} and σ_{\forall} are the existential and universal subassignments of σ . Intuitively, B_i maps to 0 iff the universal player plays according to the strategy S up to the variable u_i . Also, let B_0 be the trivial Boolean function that maps the empty assignment to 0.

It is an immediate consequence of the axiomatic equivalence of P that, for any Boolean function B , there exists a set of lines in \mathcal{L}_{P} whose conjunction has B as its associated Boolean function. ($\mathcal{A}_{\mathsf{P}}(\phi_B)$, where ϕ_B is a propositional formula representing B , is one such set.) In particular, for each $i = 0, 1, \dots, n$, there exists a set of lines \mathcal{L}_i whose conjunction has B_i as its associated Boolean function. Observe that \mathcal{L}_i essentially encodes the statement that the universal player does *not* play according to the strategy S up to the variable u_i .

By backwards induction on $i = 0, 1, \dots, n$, we show that $\mathsf{P}+\forall\text{red}$ can derive each set \mathcal{L}_i . At the final step $i = 0$, we hence prove the lemma, since B_0 is the identically zero Boolean function on an empty set of variables, which implies that \mathcal{L}_0 is semantically equivalent to \perp .

For the base case, observe that ϕ semantically entails \mathcal{L}_n , since \mathcal{L}_n encodes the statement that the universal player does not play according the whole strategy S . Therefore in a $\mathsf{P}+\forall\text{red}$ derivation from Φ one may derive each line in \mathcal{L}_n , by the axiomatic equivalence and inferential equivalence of P .

For the inductive step, let $i \in [n]$. Since u_i is the rightmost variable appearing in L_i , by universal reduction one may derive both sets of lines

$$\mathcal{L}_i^0 := \{L[u_i \mapsto 0] : L \in \mathcal{L}_i\} \quad \text{and} \quad \mathcal{L}_i^1 := \{L[u_i \mapsto 1] : L \in \mathcal{L}_i\}.$$

It is readily verified that the union of \mathcal{L}_i^0 and \mathcal{L}_i^0 semantically entails \mathcal{L}_{i-1} , which may then be derived by the inferential equivalence of P . This completes the inductive step, and the proof. \square

Soundness of the $\mathsf{P}+\forall\text{red}$ can in fact be proved for a relaxed definition of refutation in which any logically correct propositional inference is allowed, and lines introduced by universal reduction need only be semantically equivalent to the consequent from the original definition (Figure 3). A sequence satisfying these relaxed conditions we call a *semantic $\mathsf{P}+\forall\text{red}$ refutation*.

Definition 4.5 (semantic $\mathsf{P}+\forall\text{red}$ refutation). Let P be a base system. A *semantic $\mathsf{P}+\forall\text{red}$ refutation* of a QBF $\Phi = \mathcal{Q} \cdot \phi$ is a sequence $\pi = L_1, \dots, L_m$ of lines from \mathcal{L}_{P} , in which $L_m = \perp$ and each L_i satisfies at least one of the following:

- *Axiom.* $L_i \in \mathcal{A}_{\mathsf{P}}(\phi)$
- *Semantic consequence.* $\{L_1, \dots, L_{i-1}\}$ semantically entails L_i .
- *Semantic universal reduction.* $L_i \equiv L_j[\beta]$, where $j < i$, and β is a partial assignment to a universal block U of \mathcal{Q} that is rightmost in $\text{vars}(L_j)$.

Semantic refutations feature in the following section, where they are used to show that the Size-Cost-Capacity Theorem returns a genuine refutation-size lower bound. We introduce them at this point because the soundness of semantic refutations, proved in the following lemma, is used in the proof of that theorem.

Lemma 4.6. *Let P be a base system. If a QBF has a semantic $\mathsf{P}+\forall\text{red}$ refutation, then it is false.*

Proof. Let L_1, \dots, L_m be a semantic $\mathsf{P}+\forall\text{red}$ refutation of a QBF $\Phi = \mathcal{Q} \cdot \phi$. Aiming for contradiction suppose that Φ is true, and let S be a winning \exists -strategy for Φ . For each line L_i , we say that S *models* L_i if $L_i[\beta \cup S(\beta)] = \top$ for each $\beta \in \langle \text{vars}_{\forall}(\Phi) \rangle$. By induction on $i \in [m]$, we show that S models L_i . At the final step $i = n$, we therefore reach a contradiction since $L_m = \perp$ cannot be modelled by any \exists -strategy.

Since L_1 is in $\mathcal{A}_{\mathsf{P}}(\phi)$, the base case $i = 1$ follows immediately from the axiomatic equivalence of P . For the inductive step, let $i \geq 2$ and suppose that S models each L_1, \dots, L_{i-1} . The case where L_i is an axiom is identical to the base case. If L_i was derived by semantic consequence, then $L_1 \wedge \dots \wedge L_{i-1}$ semantically entails L_i , and it is easy to see that S models L_i .

It remains to consider the case where L_i was derived by semantic universal reduction. Then $L_i \equiv L_j[\beta_U]$, where $j < i$, and β_U is a partial assignment to a universal block U of \mathcal{Q} that is rightmost in $\text{vars}(L_j)$. Let $\beta \in \langle \text{vars}_{\forall}(\Phi) \rangle$, and let β' be the assignment obtained from β by overwriting the assignments to $\text{vars}(\beta_U)$ with those of β_U ; that is,

$$\beta' := (\beta \setminus \{l \in \beta : \text{var}(l) \in \text{vars}(\beta_U)\}) \cup \beta_U.$$

By the inductive hypothesis, $L_j[\beta' \cup S(\beta')] = \top$. Since $\beta_U \subseteq \beta'$, we must have $L_j[\beta_U][\beta' \cup S(\beta')] = \top$, by the restrictive closure of P . Therefore $L_i[\beta' \cup S(\beta')] = \top$. As β and β' agree on all universal blocks left of U with respect to \mathcal{Q} , $S(\beta)$ and $S(\beta')$ agree on all existential blocks left of U , by definition of \exists -strategy. Then, as U is rightmost in L_j , $\beta \cup S(\beta)$ and $\beta' \cup S(\beta')$ agree on $\text{vars}(L_i)$. Therefore $L_i[\beta \cup S(\beta)] = \top$. Thus S models L_i , completing the inductive step and the proof. \square

We arrive at the following result.

Theorem 4.7. *If P is a base system, then $\mathsf{P}+\forall\text{red}$ is a sound and complete QBF proof system.*

Proof. Immediate from Lemmata 4.4 and 4.6, and the fact that any $\mathsf{P}+\forall\text{red}$ refutation is also a semantic $\mathsf{P}+\forall\text{red}$ refutation. \square

5. THE SIZE-COST-CAPACITY THEOREM

In this section, we define the measures ‘cost’ (Subsection 5.1) and ‘capacity’ (Subsection 5.2), then state and prove our central result, the Size-Cost-Capacity Theorem (Subsection 5.3). In Subsection 5.3, we also formalise strategy extraction in $\mathsf{P}+\forall\text{red}$ and prove its correctness. Throughout this section we assume that P is a base system.

5.1. Cost. Recall that in Section 3 we proved the hardness of the equality formulas $EQ(n)$ in QU-Res by appealing to the minimum cardinality of the range of a winning strategy. For $EQ(n)$, the minimum cardinality is easy to compute because the winning strategy per instance is unique, and must contain all possible 2^n responses for the single universal block.

In order to generalise that proof method to arbitrary instances in $\mathsf{P}+\forall\text{red}$, we require a more sophisticated measure that accounts for the multiple responses collected during round-based strategy extraction in general. Fit for this purpose, we define a measure called *cost*. The cost of a false QBF is the minimum, over all winning strategies, of the largest number of responses for a single universal block.

Definition 5.1 (cost). Let $\Phi = \forall U_1 \exists E_1 \cdots \forall U_n \exists E_n \cdot \phi$ be a false QBF. Further, for each winning \forall -strategy S for Φ and each $i \in [n]$, let S_i be the function that maps each $\alpha \in \langle \text{vars}_{\exists}(\Phi) \rangle$ to the projection of $S(\alpha)$ to U_i , and let $\text{cost}(S) = \max\{|\text{rng}(S_i)| : i \in [n]\}$. The *cost* of Φ is

$$\text{cost}(\Phi) = \min\{\text{cost}(S) : S \text{ is a winning } \forall\text{-strategy for } \Phi\}.$$

The cost of $EQ(n)$ is simple to compute. There is only one winning strategy and only one universal block exhibiting 2^n responses; hence $\text{cost}(EQ(n)) = 2^n$.

Proposition 5.2. *The cost of the n^{th} equality formula is 2^n .*

5.2. Capacity. The principal notion of strategy extraction is that the \forall -player’s response (for any given round) can be read off from a suitable restriction of the refutation; as we described in Section 3, in QU-Res the response can be determined from the penultimate clause. In the general setting of $\mathsf{P}+\forall\text{red}$, we seek a method of determining the response that does not depend upon the particulars of P . We introduce the concept of a *response map* for this purpose.

For base system P and a line L in $\mathcal{L}_{\mathsf{P}}^{\mathcal{Q}}$, the *rightmost block* of L is the rightmost block Z of \mathcal{Q} for which $\text{vars}(L)$ contains a variable in Z .

Definition 5.3 (response map). Let P be a base system, let L be a line in $\mathcal{L}_{\mathsf{P}}^{\mathcal{Q}}$ whose rightmost block U is universal, and let $X := \text{vars}(L) \setminus U$. A *response map* for L with respect to \mathcal{Q} is a function $\mathcal{R} : \langle X \rangle \rightarrow \langle U \rangle$ satisfying the following for each $\alpha \in \langle X \rangle$:

If $L[\alpha]$ is not a tautology, then $\mathcal{R}(\alpha)$ falsifies $L[\alpha]$.

We call a line $L \in \mathcal{L}_P^Q$ *reducible* if its rightmost block is universally quantified.

Definition 5.4 (response map set). Given a semantic $P+\forall$ red refutation π of a QBF $Q \cdot \phi$ whose reducible lines are L_1, \dots, L_k , a *response map set* for π is a set $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ in which each \mathcal{R}_i is a response map for L_i with respect to Q .

Example 5.5. We can illustrate Definitions 5.3 and 5.4 with an example. Let Φ be a QBF, and let π be a QU-Res refutation of Φ whose reducible clauses are C_1, \dots, C_k . Consider some particular reducible clause C_i in π whose rightmost block is U_i , and let $X_i = \text{vars}(C_i) \setminus U_i$. Further, let β_i be any total assignment to U_i that falsifies the U_i -literals of C_i . The function \mathcal{R}_i mapping each $\alpha_i \in \langle X_i \rangle$ to β_i is a response map for C_i , and $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ is a response map set for π . To see this, it is sufficient to observe that, for $\alpha_i \in \langle X_i \rangle$, if α_i satisfies C_i , then $C_i[\alpha_i]$ is a tautology; otherwise, β_i falsifies $C[\alpha_i]$.

This example demonstrates that the lines in resolution – clauses – admit very simple response maps that are, in fact, constant functions. Indeed, given a QU-Res refutation we can construct a response map set which needs only one response per reducible line. In the sequel, given a refutation we shall be interested in response map sets needing the fewest responses; formally, response maps whose ranges have minimal cardinality. This minimal cardinality is determined by the *capacity* of a refutation.

Definition 5.6 (capacity). Let π be a $P+\forall$ red refutation of a QCNF $Q \cdot \phi$. The *capacity* of a response map set $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ for π is $\max_{i \in [k]} \{|\text{rng}(\mathcal{R}_i)|\}$. The *capacity* of π is the minimal capacity of any response map set for π .

Intuitively, this definition ensures that, given any $P+\forall$ red refutation π , there is some response map set for π in which each line uses at most $\text{capacity}(\pi)$ responses. Since QU-Res refutations have response map sets in which each line uses exactly one response, the capacity of any QU-Res refutation is 1.

Proposition 5.7. *Every QU-Res refutation has capacity equal to 1.*

Proof. Let π be a QU-Res refutation. The response map $\{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ given in Example 5.5 satisfies $|\text{rng}(\mathcal{R}_i)| = 1$ for each i . The capacity of any response map is at least 1; hence $\text{capacity}(\pi) = 1$. \square

5.3. Strategy extraction and Size-Cost-Capacity. We begin this subsection with a definition of strategy extraction for $P+\forall$ red. Given a semantic $P+\forall$ red refutation and response map for it, we may obtain a winning \forall -strategy by round-by-round restriction of the refutation, whereby the universal response for a given round is obtained by querying the response map on the first suitable reducible line.

Definition 5.8 (round-based strategy extraction). Let π be a semantic $P+\forall$ red refutation of a QBF $\Phi = \exists E_1 \forall U_1 \dots \exists E_n \forall U_n \cdot \phi$ with reducible lines L_1, \dots, L_k and response map set $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_k\}$. Further, for each $\alpha \in \langle \text{vars}_{\exists}(\Phi) \rangle$ and $i \in [n]$:

- let σ_0^α be the empty assignment;
- let α_i be the projection of α to E_i ;
- let F_i^α be the first line in π such that $\text{vars}(F_i^\alpha) \subseteq \bigcup_{j=1}^i (E_j \cup U_j)$ and $F_i^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i]$ is non-tautological;

- if F_i^α is some $L_{j_i^\alpha}$, then let $\beta_i^\alpha := \mathcal{R}_{j_i^\alpha}(\sigma_{i-1}^\alpha \cup \alpha_i)$, otherwise let $\beta_i^\alpha := \mathcal{R}_{j_i^U}(\sigma_{i-1}^\alpha \cup \alpha_i)$, where j_i^U is the minimal index for which U_i is the rightmost block of $L_{j_i^U}$.
- let $\sigma_i^\alpha := \sigma_{i-1}^\alpha \cup \alpha_i \cup \beta_i^\alpha$.

The *extracted strategy* for π with respect to \mathcal{R} is

$$\begin{aligned} S_{\mathcal{R}}(\pi) : \langle \text{vars}_{\exists}(\Phi) \rangle &\rightarrow \langle \text{vars}_{\forall}(\Phi) \rangle \\ \alpha &\mapsto \bigcup_{i=1}^n \beta_i^\alpha. \end{aligned}$$

It is clear that the extracted strategy is well-defined, since each F_i^α exists; in particular, for any i and α , $\text{vars}(\perp) \subseteq \bigcup_{j=1}^i (E_j \cup U_j)$ and $\perp[\sigma_{i-1}^\alpha \cup \alpha_i]$ is non-tautological, where \perp is the final line in π . (Moreover, each natural number j_i^U exists, since we can assume without loss of generality that each universal variable appears in some axiom of the refutation.) To assist the proof that the extracted strategy $S_{\mathcal{R}}(\pi)$ is indeed winning, we first prove two useful propositions which demonstrate that semantic $\text{P}+\forall\text{red}$ refutations are closed under certain restrictions. These two lemmata are essentially the $\text{P}+\forall\text{red}$ analogues of Propositions 3.2 and 3.3.

Proposition 5.9. *Let π be a semantic $\text{P}+\forall\text{red}$ refutation of a QBF Φ and let α be a partial assignment to $\text{vars}_{\exists}(\Phi)$. Then $\pi[\alpha]$ is a semantic $\text{P}+\forall\text{red}$ refutation of $\Phi[\alpha]$.*

Proof. Let $\pi = L_1, \dots, L_m$. Each L_i is derived either as an axiom, as a semantic consequence of preceding lines, or from a preceding line by semantic universal reduction. If L_i is an axiom, then $L_i[\alpha]$ may be derived as an axiom from $\Phi[\alpha]$. If L_i is a semantic consequence of the preceding lines L_{i_1}, \dots, L_{i_k} , then $L_i[\alpha]$ is a semantic consequence of the preceding lines $L_{i_1}[\alpha], \dots, L_{i_k}[\alpha]$, by the restrictive closure of P . If L_i was derived from a preceding line L_j by semantic universal reduction, then $L_i \equiv L_j[\beta]$, where β is a universal assignment. Since $\text{vars}(\alpha)$ and $\text{vars}(\beta)$ are disjoint, $L_i[\alpha]$ is semantically equivalent to $L_j[\beta][\alpha]$, and hence also to $L_j[\alpha][\beta]$ (again, by the restrictive closure of P), and may therefore be derived from $L_j[\alpha]$ by semantic universal reduction. It follows that $\pi[\alpha] = L_1[\alpha], \dots, L_m[\alpha]$ is a semantic refutation of $\Phi[\alpha]$. \square

Proposition 5.10. *Let π be a semantic $\text{P}+\forall\text{red}$ refutation of a QBF Φ whose first block U is universal, let L be the first non-tautological line appearing in π for which $\text{vars}(L) \subseteq U$, and let $\beta \in \langle U \rangle$. If $L[\beta] = \perp$ then $\pi[\beta]$ is a semantic $\text{P}+\forall\text{red}$ refutation of $\Phi[\beta]$.*

Proof. Let $\pi = L_1, \dots, L_m$, and suppose that $L[\beta] = \perp$. As in the proof of the preceding proposition, if L_i is an axiom or a semantic consequence of the preceding lines L_{i_1}, \dots, L_{i_k} , then $L_i[\beta]$ may be derived as an axiom from $\Phi[\beta]$ or as a semantic consequence of the preceding lines $L_{i_1}[\beta], \dots, L_{i_k}[\beta]$. If L_i was derived from a preceding line L_j by semantic universal reduction, then $L_i \equiv L_j[\beta']$, where β' is a partial assignment to some universal block U' . We consider two cases:

Case 1. Suppose that $U' = U$. Then L_j contains no variables right of U . Moreover, L_i cannot be L , since this would imply that L was derived by semantic universal reduction from a tautology, and is hence itself a tautology by restrictive closure of P . Therefore L_i occurs either before or after L in π . If L_i appears before L , then L_i is a tautology, in which case $L_i[\beta]$ is also a tautology (by restrictive closure of P) and can be derived as a semantic consequence from any set of lines. Otherwise L_i appears after L , and then $L_i[\beta]$ can be derived as a semantic consequence of $L[\beta] = \perp$.

Case 2. Suppose instead that $U' \neq U$. Then $\text{vars}(\beta)$ and $\text{vars}(\beta')$ are disjoint. As in the proof of Proposition 5.9, it follows that $L_i[\beta]$ can be derived by semantic universal reduction from $L_j[\beta]$.

It follows that $\pi[\beta] = L_1[\beta], \dots, L_m[\beta]$ is a semantic refutation of $\Phi[\beta]$. \square

The preceding propositions allow us to prove fairly easily that the extracted strategy is indeed winning, regardless of the choice of response map. This is the subject of the following lemma.

Lemma 5.11. *Let \mathcal{R} be a response map set for a semantic $\text{P}+\forall\text{red}$ refutation π of a QBF Φ . Then the extracted strategy for π with respect to \mathcal{R} is a winning \forall -strategy for Φ .*

Proof. Let $\Phi = \exists E_1 \forall U_1 \dots \exists E_n \forall U_n \cdot \phi$, and let S be the extracted strategy for π with respect to \mathcal{R} . We need to show that, for each $\alpha \in \langle \text{vars}_{\exists}(\Phi) \rangle$, $\alpha \cup S(\alpha)$ falsifies ϕ ; that is, $\phi[\sigma_n^\alpha] = \perp$. In fact, we prove by induction on $i = 0, 1, \dots, n$ that $\pi[\sigma_i^\alpha]$ is a semantic $\text{P}+\forall\text{red}$ refutation of $\Phi[\sigma_i^\alpha]$. Hence at the final step $i = n$, by the soundness of $\text{P}+\forall\text{red}$ (Lemma 4.6) we prove that $\Phi[\sigma_n^\alpha]$ is false, and this indeed implies $\phi[\sigma_n^\alpha] = \perp$ since σ_n^α is a total assignment to $\text{vars}(\Phi)$.

Let $\alpha \in \langle \text{vars}_{\exists}(\Phi) \rangle$. As σ_0^α is the empty assignment, the base case $i = 0$ is established from the Lemma statement. For the inductive step, let $i \in [n]$ and suppose that $\pi[\sigma_{i-1}^\alpha]$ is a semantic $\text{P}+\forall\text{red}$ refutation of $\Phi[\sigma_{i-1}^\alpha]$. Since α_i is a partial assignment to $\text{vars}_{\exists}(\Phi)$, $\pi[\sigma_{i-1}^\alpha \cup \alpha_i]$ is a semantic $\text{P}+\forall\text{red}$ refutation of $\Phi[\sigma_{i-1}^\alpha \cup \alpha_i]$, by Proposition 5.9. Now, let F_i^α be the first line in π such that $\text{vars}(F_i^\alpha) \subseteq \bigcup_{j=1}^i (E_j \cup U_j)$ and $F_i^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i]$ is non-tautological. Observe that $F_i^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i]$ is the first non-tautological line in $\pi[\sigma_{i-1}^\alpha \cup \alpha_i]$ whose variables are contained in U_i . If $\text{vars}(F_i^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i])$ contains some variable in U_i , then F_i^α is indeed some reduction line $L_{j_i}^\alpha$, and $\beta_i^\alpha \in \langle U_i \rangle$ falsifies $L_{j_i}^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i]$ by definition of response map. Otherwise $\text{vars}(F_i^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i])$ is empty, and β_i^α falsifies $L_{j_i}^\alpha[\sigma_{i-1}^\alpha \cup \alpha_i]$ trivially. In either case, $\pi[\sigma_{i-1}^\alpha \cup \alpha_i \cup \beta_i^\alpha]$ is a semantic $\text{P}+\forall\text{red}$ refutation of $\Phi[\sigma_{i-1}^\alpha \cup \alpha_i \cup \beta_i^\alpha]$, by Proposition 5.10. As $\sigma_{i-1}^\alpha \cup \alpha_i \cup \beta_i^\alpha = \sigma_i^\alpha$, this completes the inductive step, and the proof. \square

We defined round-based strategy extraction relative to an arbitrary response map. Hence, we may select one with minimum capacity; that is, we may select one that minimises the maximum number of responses extracted from a single line. By selecting such a minimal response map, we will therefore limit the capacity for lines in the refutation to contribute multiple responses to the extracted strategy. In this way we obtain our main result, which relates size, cost and capacity with a refutation-size lower bound.

Theorem 5.12. *Let π be a semantic $\text{P}+\forall\text{red}$ refutation of a QBF Φ . Then*

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

Proof. Let $\Phi = \exists E_1 \forall U_1 \dots \exists E_n \forall U_n \cdot \phi$. Further, let L_1, \dots, L_k be the reducible lines of π and let $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_k\}$ be a response map set for π with the following property: for each $i \in [k]$, the cardinality of the range of \mathcal{R}_i is minimal; that is, for each response map \mathcal{R}'_i for L_i , $|\text{rng}(\mathcal{R}_i)| \leq |\text{rng}(\mathcal{R}'_i)|$. Further, let S be the extracted strategy for π with respect to \mathcal{R} . By Lemma 5.11, S is a winning \forall -strategy for Φ . By definition of cost, there exists some universal block U_i such that $|\{\text{proj}(S(\alpha), U_i) : \alpha \in \langle \text{vars}_{\exists}(\Phi) \rangle\}|$ is at least $\text{cost}(\Phi)$. Equivalently, there exists some $i \in [n]$ such that $|\{\beta_i^\alpha : \alpha \in \langle \text{vars}_{\exists}(\Phi) \rangle\}| \geq \text{cost}(\Phi)$. Each β_i^α is in the range of the response map $\mathcal{R}_{j_i}^\alpha$ for some reducible line $L_{j_i}^\alpha$ in π . By definition of

capacity, the cardinality of the range of $\mathcal{R}_{i,\alpha}$ is at most $\text{capacity}(\pi)$. Hence we must have $|\pi| \geq \text{cost}(\Phi)/\text{capacity}(\pi)$, for otherwise the number of lines in π would exceed its size. \square

We emphasize that the lower bound given by Theorem 5.12 applies to *semantic* $\text{P}+\forall\text{red}$ refutations. Therefore, any result obtained as an application of our technique is by definition a genuine lower bound.

Since QU-Res derivations have unit capacity (Proposition 5.7), the Size-Cost-Capacity Theorem tells us that *cost alone* is an absolute refutation-size lower bound in that system.

Corollary 5.13. *Let π be a QU-Res refutation of a QBF Φ . Then $|\pi| \geq \text{cost}(\Phi)$.*

For a straightforward application of our technique, we could use Corollary 5.13 to prove the hardness of the equality formulas in QU-Res (Theorem 3.5), as a direct consequence of their exponential cost (Proposition 5.2), as opposed to the direct method of Section 3.

6. CAPACITY BOUNDS

In this section, we demonstrate that this lower bound technique is widely applicable to QBF proof systems by showing upper bounds on the capacity of proofs in the QBF versions of two commonly studied propositional proof systems: Cutting Planes (Subsection 6.1) and Polynomial Calculus with Resolution (Subsection 6.2). These proof systems represent two distinct approaches to propositional proof systems, via integer linear programming and algebraic methods respectively. Both proof systems are known to p-simulate resolution, and similarly the QBF proof systems obtained with the addition of the \forall -reduction rule both p-simulate QU-Res. Our capacity upper bound for PCR+ $\forall\text{red}$ proofs is particularly noteworthy as it is not constant, but depends on the size of the proof. We conclude this section with an example of the limits of this technique, namely a Frege+ $\forall\text{red}$ proof with large capacity.

6.1. Cutting Planes. The first proof system we analyse is Cutting Planes [CCT87] and its extension to QBFs, CP+ $\forall\text{red}$ [BCMS18b]. Inspired by integer linear programming, Cutting Planes translates a CNF into an equivalent system of linear inequalities, and from these derives the trivial falsehood $0 \geq 1$. Replacing the axiom rules by any unsatisfiable set of inequalities, Cutting Planes is in fact complete for any set of linear inequalities without integer solutions, and is implicationally complete with the inclusion of the trivial truth $0 \geq -1$ as an axiom. Cutting Planes is therefore a base system. As we focus only on Cutting Planes as a proof system for unsatisfiable CNFs, for which the addition of $0 \geq -1$ has no effect on proof size, our definition of Cutting Planes omits this axiom, in keeping with previous work.

Cutting Planes has two inference rules: the linear combination rule and the division rule. The linear combination rule infers from two inequalities, some linear combination of these inequalities with non-negative integer coefficients. The division rule allows division by any integer $c > 0$ if c divides the coefficient of each variable; note that c need not divide the constant term in the inequality.

Definition 6.1 (Cutting Planes [CCT87]). A line L in a Cutting Planes CP proof is a linear inequality $a_1x_1 + \dots + a_nx_n \geq A$ where $\text{vars}(L) = \{x_1, \dots, x_n\}$ and $a_1, \dots, a_n, A \in \mathbb{Z}$.¹⁰

¹⁰For convenience and clarity, we may refer to lines in \mathcal{L}_{CP} and $\mathcal{L}_{\text{CP}}^{\forall}$ using equivalent linear inequalities not in this precise form. Similarly, the result of any \forall -reduction is expressed as a line of this form.

A Cutting Planes derivation of a line $L \in \mathcal{L}_{\text{CP}}$ from a CNF ϕ consists of a sequence of lines L_1, \dots, L_m where $L_m = L$ and each line $L_i \in \mathcal{L}_{\text{CP}}$ is either an instance of an axiom rule, or is derived from the previous lines by an inference rule (Figure 4). A CP refutation of ϕ is a derivation of $0 \geq 1$ from ϕ .

Clause Axiom:	$\overline{\sum_{l \in C} R(l) \geq 1}$	for any clause $C \in \phi$, where $R(x) = x$, $R(-x) = 1 - x$
Boolean Axiom:	$\overline{x \geq 0} \quad \overline{-x \geq -1}$	for any variable x
Linear combination:	$\frac{\sum_i a_i x_i \geq A \quad \sum_i b_i x_i \geq B}{\sum_i (\alpha a_i + \beta b_i) x_i \geq \alpha A + \beta B}$	for any $\alpha, \beta \in \mathbb{N}$
Division:	$\frac{\sum_i c a_i x_i \geq A}{\sum_i a_i x_i \geq \lceil \frac{A}{c} \rceil}$	for any non-zero $c \in \mathbb{N}$

Figure 4: The rules of the Cutting Planes proof system (CP)

It is straightforward to see that Cutting Planes p -simulates resolution.¹¹ Indeed, it is strictly stronger than resolution, as there are short CP proofs of the pigeonhole principle formulas, which are known to require large proofs in resolution [Hak85]. The same argument shows that CP+ \forall red is exponentially stronger than QU-Res.

Despite the apparent strength of CP+ \forall red compared to QU-Res, any proof in CP+ \forall red still only has unit capacity. This comes about as the left hand side of any inequality L is simply a linear combination of variables. We can therefore construct a constant response map defined by the coefficients of the variables of the rightmost universal block, with the aim of minimising the left hand side of the inequality.

Proposition 6.2. *For every CP+ \forall red derivation π , $\text{capacity}(\pi) = 1$.*

Proof. Consider any $L \in \mathcal{L}_{\text{CP}}^{\mathcal{Q}}$ with rightmost block U universally quantified, and with $X = \text{vars}(L) \setminus U$. Then L is of the form $\sum_{x \in X} a_x x + \sum_{u \in U} b_u u \geq c$, where $a_x, b_u, c \in \mathbb{Z}$ for all $x \in X, u \in U$.

Define the assignment $\beta_L \in \langle U \rangle$ by assigning each variable u to 0 if $b_u \geq 0$, and to 1 otherwise. It is clear that this response minimises the value of $\sum_{u \in U} b_u u$, and thus minimises the left hand side of the inequality in the line $L[\alpha]$ for any $\alpha \in \langle X \rangle$. If any response falsifies $L[\alpha]$, a response which minimises this term must do so, so the map $\mathcal{R}_L : \langle X \rangle \rightarrow \langle U \rangle$ defined by $\mathcal{R}_L(\alpha) = \beta_L$ is a response map for L .

For any CP+ \forall red proof π , the set $\{\mathcal{R}_L : L \text{ is a reducible line of } \pi\}$ is a response map set for π , witnessing that $\text{capacity}(\pi) = 1$. \square

Having established that any CP+ \forall red proof has capacity 1, we can apply the Size-Cost-Capacity Theorem (Theorem 5.12). This yields a lower bound on the size of CP+ \forall red proofs of a QBF Φ determined solely by the cost of Φ .

¹¹A proof system P p -simulates a system Q if each Q -proof can be transformed in polynomial time into a P -proof of the same formula [CR79].

Corollary 6.3. *Let π be a $\text{CP}+\forall\text{red}$ refutation of a QBF Φ . Then $|\pi| \geq \text{cost}(\Phi)$.*

Hence, even in the stronger system of $\text{CP}+\forall\text{red}$, we still have a straightforward proof that refutations of the equality formulas require size 2^n by looking at the cost of the formulas and using Size-Cost-Capacity.

6.2. Polynomial Calculus. Polynomial Calculus [CEI96] presents an algebraic approach to proving unsatisfiability. A CNF ϕ is translated into a set of polynomials over an arbitrary field \mathbb{F} , such that any assignment where all polynomials evaluate to zero corresponds to a satisfying assignment for ϕ , and vice versa. The field \mathbb{F} is often assumed to be \mathbb{Q} , but since all axioms and any solutions use only the constants 0 and 1, Polynomial Calculus (as a proof system for CNFs) can be defined over any field \mathbb{F} .

Formally, Polynomial Calculus works with polynomial equations where the right hand side is 0. A Polynomial Calculus refutation of a set of polynomials is a derivation of the equation $1 = 0$, which is enough to show that the set of polynomials has no common solution. The inference rules permit deriving any linear combination of two previous lines, or multiplying any line by a single variable.

As a propositional proof system, the axioms of Polynomial Calculus are polynomials equivalent to each clause in a CNF. Given a clause C in a CNF, the corresponding polynomial axiom is $\prod_{l \in C} V(l) = 0$, where $V(x) = x$ and $V(\neg x) = (1 - x)$.¹² We also include the axioms $x^2 - x = 0$ to ensure only Boolean solutions are possible. With these Boolean axioms added, Polynomial Calculus is implicationaly complete [CEI96, FLM⁺13], ensuring that Polynomial Calculus is a base system.

Proof size in Polynomial Calculus is measured by the number of *monomials* in the lines of the proof. By this measure of proof size, Polynomial Calculus cannot even simulate resolution, as the clause $\neg x_1 \vee \dots \vee \neg x_n$ would translate to $(1 - x_1) \dots (1 - x_n) = 0$, which contains 2^n monomials. As a result of this issue, a modification of Polynomial Calculus was introduced in [ABSRW02], using variables \bar{x} representing $\neg x$. The inference rules remain the same, but we add the axioms $x + \bar{x} - 1 = 0$ for each $x \in \text{vars}(\phi)$ to ensure that x and \bar{x} take opposite values.

Definition 6.4 (Polynomial Calculus with Resolution [ABSRW02]). Fix an arbitrary field \mathbb{F} . Given a CNF ϕ , lines in a PCR derivation from ϕ are polynomials in the variables $\{x, \bar{x} : x \in \text{vars}(\phi)\}$. A PCR derivation of a line $L \in \mathcal{L}_{\text{PCR}}$ from a CNF ϕ is a sequence of lines L_1, \dots, L_m in \mathcal{L}_{PCR} such that $L_m = L$, and each L_i is an instance of an axiom rule, or derived from previous lines by one of the inference rules (Figure 5). A PCR refutation of ϕ is a derivation of the line $1 = 0$.

As Polynomial Calculus with Resolution is clearly at least as strong as Polynomial Calculus (and in fact is strictly stronger), we focus here only on the version with Resolution. The capacity upper bounds shown, and consequent proof size lower bounds, all hold for Polynomial Calculus as well.

In contrast to QU-Res and $\text{CP}+\forall\text{red}$, not all proofs in $\text{PCR}+\forall\text{red}$ have unit capacity. For a simple example, consider the line $x(1 - u) + (1 - x)u = 0$. This polynomial clearly evaluates to 0 if and only if $x = u$, so the only winning response for u is to play $u = 1 - x$.

¹²Assigning the algebraic variable x to 0 is therefore equivalent to assigning the corresponding Boolean variable to 1, and vice versa. This ‘swapping’ of truth values is a common convention for algebraic proof systems.

Clause Axiom:	$\overline{\prod_{l \in C} V(l) = 0}$	for any clause $C \in \phi$, where $V(x) = x$, $V(\neg x) = \bar{x}$
Boolean Axiom:	$\overline{y^2 - y = 0} \quad \overline{x + \bar{x} - 1 = 0}$	for any $y \in \{x, \bar{x} : x \in \text{vars}(\phi)\}$ for any $x \in \text{vars}(\phi)$
Linear combination:	$\frac{p(\vec{x}) = 0 \quad q(\vec{x}) = 0}{\alpha \cdot p(\vec{x}) + \beta \cdot q(\vec{x}) = 0}$	for any $\alpha, \beta \in \mathbb{F}$
Multiplication:	$\frac{p(\vec{x}) = 0}{y \cdot p(\vec{x}) = 0}$	for any $y \in \{x, \bar{x} : x \in \text{vars}(\phi)\}$

Figure 5: The rules of Polynomial Calculus with Resolution (PCR)

The unique response map \mathcal{R} for this line therefore has $\text{rng}(\mathcal{R}) = 2$, so if a PCR+ \forall red proof π contains such a line then $\text{capacity}(\pi) \geq 2$.

Given a QBF Φ with n variables in the universal block U , it is possible to construct a line in $\mathcal{L}_{\text{PCR}}^{\mathcal{Q}}$ which requires a response map with an exponential size range. However, as previously mentioned, the size of a proof in PCR+ \forall red is not measured by the number of lines in the proof, but by the number of monomials in the proof. To provide a suitable capacity upper bound, it is thus sufficient to upper bound the size of a response set for a line by the number of monomials in that line. That is, we show that any proof with large capacity must contain a large line, and so the proof itself is large.

Proposition 6.5. *If π is a PCR+ \forall red proof where each line contains at most M monomials, then $\text{capacity}(\pi) \leq M$.*

The key observation used in proving this bound is that, rather than considering the specific assignment to the universal variables, it is sufficient to consider the value the response gives to the monomials of a line $L \in \mathcal{L}_{\text{PCR}}^{\mathcal{Q}}$. If two assignments $\beta, \beta' \in \langle U \rangle$ set the same monomials to zero and retain the same monomials, only one of these assignments need be in the range of a response map as $L[\beta] = L[\beta']$ and so β and β' falsify $L[\alpha]$ for precisely the same $\alpha \in \langle X \rangle$.

We therefore consider the set of different assignments to these monomials needed, as a subset of $\{0, 1\}^M$. We construct a response map for L such that the responses, when considered in this way, form a linearly independent subset of $\{0, 1\}^M$, providing the upper bound.

Proof of Proposition 6.5. To upper bound $\text{capacity}(\pi)$, it is enough to find, for any reducible line $L \in \pi$, a response map \mathcal{R}_L such that $|\text{rng}(\mathcal{R}_L)| \leq M$. To this end, fix such a line $L \in \pi$ with at most M monomials. Let U to be the rightmost universal block present in L , and $X := \text{vars}(L) \setminus U$, and write L as

$$\sum_{j=1}^N f_j v_j = 0$$

where each f_j is a polynomial (*not* necessarily a single monomial) in the variables of X and the v_j are distinct monomials in the variables U . Denote by $f_j[\alpha]$ and $v_j[\beta]$ the values (in \mathbb{F}) obtained by evaluating f_j , respectively v_j , according to the assignments $\alpha \in \langle X \rangle$, respectively $\beta \in \langle U \rangle$.

Observe that since L contains at most M monomials, $N \leq M$. We now construct a response map $\mathcal{R}_L : \langle X \rangle \rightarrow \langle U \rangle$ for which $|\text{rng}(\mathcal{R}_L)| \leq N$. The response map set for π consisting of \mathcal{R}_L for every reducible line $L \in \pi$ demonstrates that $\text{capacity}(\pi) \leq M$.

First, enumerate the elements of $\langle X \rangle$ as $\langle X \rangle = \{\alpha_1, \dots, \alpha_m\}$. We construct a sequence of functions $\mathcal{R}_L^i : \{\alpha_1, \dots, \alpha_i\} \rightarrow \langle U \rangle$ such that \mathcal{R}_L^0 is the empty function, \mathcal{R}_L^i extends \mathcal{R}_L^{i-1} , and $\mathcal{R}_L^m = \mathcal{R}_L$ is a complete response map for L . Moreover, for each $0 \leq i \leq m$, $|\text{rng}(\mathcal{R}_L^i)| \leq N$, in particular $|\text{rng}(\mathcal{R}_L^m)| \leq N$.

The construction of the function \mathcal{R}_L^m involves, in effect, going through each of the assignments $\alpha \in \langle X \rangle$ in turn and choosing the response of \mathcal{R}_L^m on α . At each stage, if there is a suitable response that has been chosen before, we choose it again as we are aiming to minimise $|\text{rng}(\mathcal{R}_L)|$. If there is no suitable previously chosen response, we choose a suitable response and then show that by evaluating the monomials v_j on the responses, there is an injection from the set of chosen responses into a linearly independent subset of \mathbb{F}^N .

If L is a tautology, then for any $\alpha \in \langle X \rangle$, $L[\alpha]$ cannot be falsified, and so any constant function $\mathcal{R}_L : \langle X \rangle \rightarrow \langle U \rangle$ is a response map with range 1. We therefore assume without loss of generality that L is not a tautology, and in particular that $L[\alpha_1]$ is falsifiable.

Let $R_i := \text{rng}(\mathcal{R}_L^i)$. For any $\beta \in \langle U \rangle$, denote by $\vec{v}[\beta]$ the vector $(v_1[\beta], \dots, v_N[\beta]) \in \{0, 1\}^N$. We now define a response $\beta_i \in \langle U \rangle$ which extends \mathcal{R}_L^{i-1} to \mathcal{R}_L^i such that \mathcal{R}_L^i can be extended to a response map for L , i.e. whenever $L[\alpha_i]$ is falsifiable, $L[\alpha_i \cup \beta_i] = \perp$.

Furthermore, we show by induction on i that the set $V_i = \{\vec{v}[\beta] : \beta \in R_i\}$ is linearly independent (as a subset of \mathbb{F}^N). Since $V_i \subseteq \{0, 1\}^N$, and $|V_i| = |R_i|$, this provides the upper bound on $|R_i|$.

The empty function \mathcal{R}_L^0 ensures that $R_0 = \emptyset$, which is linearly independent. Given a function \mathcal{R}_L^{i-1} , we define \mathcal{R}_L^i as follows:

- If $L[\alpha_i]$ is a tautology, i.e. $\sum_{j=1}^N f_j[\alpha_i]v_j[\beta] = 0$ for all $\beta \in \langle U \rangle$, a response map can map α_i to any response in $\langle U \rangle$. Since $i \neq 1$, R_{i-1} is non-empty, so define $\mathcal{R}_L^i(\alpha_i) = \mathcal{R}_L^{i-1}(\alpha_1)$, whence $R_i = R_{i-1}$. As $V_i = V_{i-1}$ and V_{i-1} is linearly independent, V_i is also linearly independent.
- If $L[\alpha_i \cup \beta] = \perp$ for some $\beta \in R_{i-1}$, then define $\mathcal{R}_L^i(\alpha_i) = \beta$. Now, as previously, $R_i = R_{i-1}$ and V_i is linearly independent as $V_i = V_{i-1}$.
- Else $L[\alpha_i]$ is not a tautology, but $L[\alpha_i \cup \beta] = \top$ for all $\beta \in R_{i-1}$. Suppose $R_{i-1} = \{\beta_1, \dots, \beta_k\}$, then there exists some β_{k+1} such that $L[\alpha_i \cup \beta_{k+1}] = \perp$, since $L[\alpha_i]$ is not a tautology. Define $\mathcal{R}_L^i(\alpha_i) = \beta_{k+1}$, with $R_i = \{\beta_1, \dots, \beta_k, \beta_{k+1}\}$. We now need only show that $V_i = \{\vec{v}[\beta_j] : 1 \leq j \leq k+1\}$ is linearly independent.

For any $1 \leq l \leq k$, $L[\alpha_i][\beta_l] \neq L[\alpha_i][\beta_{k+1}]$, so $\vec{v}[\beta_{k+1}] \neq \vec{v}[\beta_l]$. Suppose there is some linear dependence relation on V_i . Since V_{i-1} is linearly independent, $\vec{v}[\beta_{k+1}]$ must have a non-zero coefficient in any such linear combination, hence there are constants $c_1, \dots, c_k \in \mathbb{F}$ such that $\sum_{t=1}^k c_t \vec{v}[\beta_t] = \vec{v}[\beta_{k+1}]$. If such constants exist, we can use the same constants to construct a linear combination of the $\sum_{j=1}^N f_j[\alpha_i]v_j[\beta_t]$, by assumption all equal to zero,

summing to $\sum_{j=1}^N f_j[\alpha_i]v_j[\beta_{k+1}]$, which by choice of β_{k+1} is non-zero.

$$0 = \sum_{t=1}^k c_t \sum_{j=1}^N f_j[\alpha_i]v_j[\beta_t] = \sum_{j=1}^N f_j[\alpha_i] \sum_{t=1}^k c_t v_j[\beta_t] = \sum_{j=1}^N f_j[\alpha_i]v_j[\beta_{k+1}] \neq 0$$

From this contradiction, we conclude that the constants c_t do not exist, and thus that V_i is a linearly independent set.

The set V_m forms a linearly independent set which is a subset of \mathbb{F}^N , so has cardinality at most N . By the construction of each \mathcal{R}_L^i , R_i cannot contain $\beta \neq \beta'$ with $\vec{v}(\beta) = \vec{v}(\beta')$, so $|R_m| = |V_m|$. Consequently, the map $\mathcal{R}_L = \mathcal{R}_L^m$ satisfies $|\text{rng}(\mathcal{R}_L)| \leq N \leq M$. We therefore obtain the response map set $\{\mathcal{R}_L : L \in \pi \text{ is reducible}\}$ for π and conclude that $\text{capacity}(\pi) \leq M$. \square

The effect of this bound is to show that PCR+ \forall red proofs with large capacity also have large size, as they must contain lines with a large number of monomials. This provides a lower bound for PCR+ \forall red proofs of a QBF Φ based solely on $\text{cost}(\Phi)$, since small proofs also have small capacity.

Corollary 6.6. *Let π be a PCR+ \forall red refutation of a QBF Φ . Then $|\pi| \geq \sqrt{\text{cost}(\Phi)}$.*

Proof. As the size of π is measured by the number of monomials, each line of π contains at most $|\pi|$ monomials, and so by Proposition 6.5, $\text{capacity}(\pi) \leq |\pi|$. Applying Size-Cost-Capacity (Theorem 5.12), we conclude that $|\pi| \geq \frac{\text{cost}(\Phi)}{|\pi|}$, i.e. $|\pi| \geq \sqrt{\text{cost}(\Phi)}$. \square

As for QU-Res and CP+ \forall red, this immediately gives a lower bound of $2^{\Omega(n)}$ for any proof of the equality formulas in PCR+ \forall red.

6.3. Proofs with large capacity. We conclude this section by noting that our technique cannot be applied to some of the more powerful proof systems. These proof systems use lines which are able to concisely express more complex Boolean functions which require large response sets. The example we give is a proof of the equality formulas in the proof system Frege+ \forall red. The Frege+ \forall red proof system is the strongest proof system we discuss in this paper, and no superpolynomial lower bounds on proof size are known in the propositional system Frege, nor in the QBF proof system Frege+ \forall red.

The Frege proof system is a ‘textbook’ propositional proof system, in which lines are arbitrary formulas in propositional variables, the constants \top, \perp and the connectives \wedge, \vee, \neg . The rules of a Frege system comprise a set of axiom schemes and inference rules, which must be implicational complete [CR79]; all such systems are equivalent [CR79, Kra95].

Proposition 6.7. *There is a Frege+ \forall red refutation of the n^{th} equality formula $EQ(n)$ with size $O(n)$.*

Proof. From the lines $x_i \vee u_i \vee \neg t_i$ and $\neg x_i \vee \neg u_i \vee \neg t_i$, there is a constant size Frege+ \forall red derivation of the line $L_i = [(x_i \vee u_i) \wedge (\neg x_i \vee \neg u_i)] \vee \neg t_i$. Successively applying the resolution rule, which Frege can simulate with constant size, to the lines L_i and the line $t_1 \vee \dots \vee t_n$ results in the line

$$L = \bigvee_{i=1}^n [(x_i \vee u_i) \wedge (\neg x_i \vee \neg u_i)].$$

Let K_m be the line $\bigvee_{i=1}^m [(x_i \vee u_i) \wedge (\neg x_i \vee \neg u_i)]$, so $K_n = L$ and $K_0 = \perp$. For any m , K_{m-1} can be derived from K_m by first \forall -reducing to obtain $K_m[u_m/0] \equiv K_{m-1} \vee x_m$ and $K_m[u_m/1] \equiv K_{m-1} \vee \neg x_m$, and then resolving these on x_m .

Deriving each K_m in turn from K_{m+1} provides a linear size refutation of L , and hence a linear size refutation of $EQ(n)$. \square

The Frege+ \forall red proof π defined in the proof above only uses formulas of depth 3 (i.e. there are at most two alternations between \wedge and \vee) and so π is also a refutation in the more restrictive systems AC^0 -Frege+ \forall red and AC_3^0 -Frege+ \forall red. By the Size-Cost-Capacity Theorem (Theorem 5.12), we know that $\text{capacity}(\pi)$ is of exponential size. We can show this directly by considering the reducible line L from π .

Any winning response for L to an assignment $\alpha \in \langle \{x_1, \dots, x_n\} \rangle$ must falsify $(x_i \vee u_i) \wedge (\neg x_i \vee \neg u_i)$ for each $1 \leq i \leq n$. The unique winning response to α is therefore to play β such that $\beta(u_i) = \alpha(x_i)$. Since there are 2^n distinct assignments in $\langle \{x_1, \dots, x_n\} \rangle$, the range of the unique response map for L is $\langle \{u_1, \dots, u_n\} \rangle$, which has size 2^n , despite L being a Frege+ \forall red line of size polynomial in n , and so $\text{capacity}(\pi) \geq 2^n$.

7. RANDOMLY GENERATED FORMULAS WITH LARGE COST

In the previous section, we saw that Size-Cost-Capacity can be used to simultaneously show lower bounds in many different QBF proof systems simply by examining the cost of QBFs. We now construct a class of randomly generated QBFs, denoted $Q(n, m, c)$ which with high probability are false and have large cost for appropriate values of m and c . By showing a lower bound on cost, we immediately obtain lower bounds on proof size in QU-Res, CP+ \forall red and PCR+ \forall red for these random QBFs.

The model used for these formulas is a simple extension of previous constructions of random QBFs, and our lower bounds have both theoretical and practical significance. On the theoretical side, the high cost of these formulas demonstrates that our lower bound technique applies to a large number of QBFs, and not just a few handcrafted instances. For practical QBF solving, it is important to have a large number of instances on which to test a solver. Having identified cost as a simple reason for QBF proof size lower bounds, the utility of generating many QBFs with large cost is clear.

Definition 7.1. 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:

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

Specifying that clauses contain a given number of literals from different sets may seem unusual, especially to readers familiar with random SAT instances, however it is widely used in the study of randomly generated QBFs [CI05, CDER15]. If any clause in the matrix of a QBF contains only literals on universal variables, then it is easy to see that the QBF is false, and that all proof systems P+ \forall red have a constant size refutation using only this clause. Specifying that all clauses must contain a given number of literals from different sets of variables avoids this issue by guaranteeing that all clauses contain existential variables. It is natural that we would also expect clauses in a QBF to contain universal variables.

We define the QBFs $\Psi_i := \exists Y_i \forall X_i \cdot \bigwedge_{j=1}^{cn} C_i^j$, which are generated using the same clauses as the (1,2)-QCNF model (Definition 7.4). We can easily see that $Q(n, m, c) \equiv \bigvee_{i=1}^n \Psi_i$ is simply a disjunction of n QBFs generated using this commonly studied model for random QBFs. This equivalence allows us to better understand $Q(n, m, c)$ by studying the randomly generated Ψ_i .

We first show that, for suitable values of m and c , with probability $1 - o(1)$, all Ψ_i are false, and furthermore for a linear number of the Ψ_i , $\text{cost}(\Psi_i) \geq 2$. By observing that a winning strategy for $Q(n, m, c)$ is a winning strategy for all the Ψ_i simultaneously, we can then show a cost lower bound for $Q(n, m, c)$.

We first focus on proving that, with suitable values for m and c , $Q(n, m, c)$ is false with high probability. This is equivalent to showing that, with high probability, each of the Ψ_i is false. In each Ψ_i , any winning assignment for the \exists -player must satisfy an existential literal in every clause. If not, there would be a winning \forall -strategy constructed by finding a clause where both existential literals were false, and setting the universal literal in that clause to false. Determining the truth of Ψ_i can therefore be reduced to determining the satisfiability of the 2-SAT problem defined by the existential parts of the clauses C_i^j . We can then use the following result on the satisfiability of random 2-SAT formulas, shown independently by Chvátal and Reed [CR92], Goerdt [Goe96] and de la Vega [dlV98], to obtain the falsity of the Ψ_i . We state it here with a tighter probability lower bound of $1 - o(n^{-1})$ proved by de la Vega in [dlV01], which is necessary for our present work.

Theorem 7.2 (de la Vega [dlV01]). *Let Φ be a random 2-SAT formula on n propositional variables containing cn clauses selected uniformly at random. If $c > 1$ then Φ is unsatisfiable with probability $1 - o(n^{-1})$.*

The following lemma is equivalent to the statement that, with the same bounds on m and c , $Q(n, m, c)$ is false with probability $1 - o(1)$. This is a fairly immediate consequence of Theorem 7.2; we need only check that it is sufficiently likely that the existential parts of the clauses of the Ψ_i satisfy the conditions of Theorem 7.2. The possibility of repeating an existential clause many times with different universal literals makes this non-trivial, but the proof is relatively straightforward.

Lemma 7.3. *For each $1 \leq i \leq n$, let ψ_i be a set of cn clauses picked uniformly at random from the set of clauses containing 1 literal from $X_i = \{x_i^1, \dots, x_i^m\}$ and 2 literals from $Y_i = \{y_i^1, \dots, y_i^n\}$. If $m \leq \log_2(n)$ and $c > 1$, then with probability $1 - o(1)$, $\Psi_i := \exists Y_i \forall X_i \cdot \psi_i$ is false for all $1 \leq i \leq n$.*

Proof. For the QBFs Ψ_i to be false, it is sufficient for the 2-SAT problem generated by taking only the existential parts of the clauses to be false, as the universal response need only respond by falsifying the universal literal on some unsatisfied existential clause. However, it is possible that clauses in ψ_i contain the same existential literals and differ only in the universal literal. In order to use Theorem 7.2, we need to show that there is some constant $k > 1$ such that, for each $i \in [n]$, the clauses of ψ_i contain at least kn distinct pairs of existential literals with high probability.

For each ψ_i , there are $4 \binom{n}{2}$ choices for the existential literals of a clause, and $2m \leq 2 \log(n)$ possible universal literals. The total number of possible clauses is therefore at most $4n(n-1) \log(n)$.

Let k be some constant with $1 < k < c$. To determine the probability of ψ_i containing at least kn distinct clauses in the existential variables, consider successively making cn random

choices of clause from the $4n(n-1)\log(n)$ possible clauses. If, on choosing a clause, fewer than kn distinct existential clauses have been chosen, the probability of a randomly chosen clause having existential part distinct from all previously chosen clauses is at least

$$\frac{4n(n-1)\log(n) - 2kn\log(n)}{4n(n-1)\log(n)} = 1 - \frac{k}{2(n-1)}.$$

Define the selection of a clause to be successful if it either selects a clause with existential part distinct from that of the previous clauses, or if kn distinct existential clauses have already been selected. The probability of any selection being successful is therefore at least $1 - \frac{k}{2(n-1)}$. It is enough to show that if we select cn clauses, with a probability $1 - \frac{k}{2(n-1)}$ of success for each selection, then the probability of fewer than kn successes is $O(e^{-n})$.

The distribution of the random variable Z , the total number of successes, is a sum of cn Bernoulli random variables with $p = 1 - \frac{k}{2(n-1)}$. Substituting these values into Hoeffding's inequality, we obtain

$$P(Z \leq kn) \leq \exp\left(-2\frac{\left(cn - \frac{kc n}{2(n-1)} - kn\right)^2}{cn}\right) = \exp\left(-\frac{2(c-k)^2}{c}n + O(1)\right)$$

and so $P(Z > kn) = 1 - \frac{1}{e^{\Omega(n)}} = 1 - o(n^{-1})$.

The probability that a given Ψ_i is false is at least the probability of it containing at least kn distinct existential clauses and the first kn distinct such clauses being unsatisfiable. Given the clauses of ψ_i were chosen uniformly at random, each set of kn existential clauses is equally likely to be chosen, so the probability these clauses are unsatisfiable is $1 - o(n^{-1})$, by Theorem 7.2. The probability of Ψ_i being false is therefore at least $P(Z > kn) \cdot (1 - o(n^{-1})) = 1 - o(n^{-1})$.

Finally, the selection of clauses for each Ψ_i is independent of clauses chosen in any other Ψ_i , and so the probability of all being false is $(1 - o(n^{-1}))^n = 1 - o(1)$. \square

It remains to show that $\text{cost}(Q(n, m, c))$ is large. Again, we first look at the cost of Ψ_i , and observe that, for $m \leq \log_2(n)$ and $1 < c < 2$, $\text{cost}(\Psi_i) \geq 2$ with probability $1 - o(1)$. Winning responses for $Q(n, m, c)$ are simultaneous winning responses for each of the Ψ_i . As many of the Ψ_i require multiple distinct responses, it is reasonable to expect that the number of responses to falsify all of them is large. With a careful choice of the parameters m and c , we can indeed force $Q(n, m, c)$ to have a large cost with high probability.

To prove $\text{cost}(\Psi_i) \geq 2$, it is only necessary to show that $\text{cost}(\Psi_i) \neq 1$, i.e. that any winning \forall -strategy $S : \langle Y_i \rangle \rightarrow \langle X_i \rangle$ for Ψ_i is not constant. If there is a constant winning \forall -strategy, say $S(\alpha) = \beta$ for all $\alpha \in \langle Y_i \rangle$, then β also constitutes a winning \forall -strategy for $\Psi'_i = \forall X_i \exists Y_i \cdot \psi_i$.

Definition 7.4 (Chen and Interian [CI05]). A (1,2)-QCNF is a QBF of the form $\forall X \exists Y \cdot \phi(X, Y)$ where $X = \{x_1, \dots, x_m\}$, $Y = \{y_1, \dots, y_n\}$ and $\phi(X, Y)$ is a 3-CNF formula in which each clause contains one universal literal and two existential literals.

If a winning \forall -strategy for Ψ'_i exists, then Ψ'_i is false. However, for $c < 2$, Ψ'_i is known to be true with high probability.

Theorem 7.5 (Creignou et al. [CDER15]). *Let Φ be a (1,2)-QCNF in which $\phi(X, Y)$ contains cn clauses picked uniformly at random from the set of all suitable clauses. If $m \leq \log_2(n)$, and if $c < 2$, then Φ is true with probability $1 - o(1)$.*

We therefore pick the parameter c to lie between the lower bound from Theorem 7.2 and the upper bound from Theorem 7.5. From these results, we see that for $1 < c < 2$, $\exists Y \forall X. \psi$ is false, but $\forall X \exists Y. \psi$ is true with high probability. Any constant winning \forall -strategy for $\exists Y \forall X. \psi$ is also a winning \forall -strategy for $\forall X \exists Y. \psi$, whence the latter is false. This gives us the bound $\text{cost}(\Psi_i) \geq 2$ with high probability.

Lemma 7.6. *Let ψ be a set of cn clauses picked uniformly at random from the set of clauses containing 1 literal from the set $X = \{x_1, \dots, x_m\}$ and 2 literals from $Y = \{y_1, \dots, y_n\}$. If $1 < c < 2$ and $m \leq \log_2(n)$, then with probability $1 - o(1)$, $\Psi := \exists Y \forall X. \psi$ is false, and $\text{cost}(\Psi) \geq 2$.*

Proof. Observe from the proof of Lemma 7.3 that, as $c > 1$, Ψ is false with probability $1 - o(n^{-1})$ and $\text{cost}(\Psi) \geq 1$.

Suppose $\text{cost}(\Psi) = 1$, then there is some $\beta \in \langle X \rangle$ such that β is a winning response for any $\alpha \in \langle Y \rangle$. That is, for any $\alpha \in \langle Y \rangle$, $\psi[\alpha][\beta] = \perp$. We can use β as a winning strategy for $\Psi' = \forall X \exists Y. \psi$, defining $S'(\emptyset) = \beta$. Since $\psi[\beta][\alpha] = \perp$ for all $\alpha \in \langle Y \rangle$, S' is a winning \forall -strategy and so Ψ' is false. However since $c < 2$, Ψ' is false with probability $o(1)$ (Theorem 7.5), and so such a $\beta \in \langle X \rangle$ exists with probability $o(1)$.

The probability that Ψ is false and $\text{cost}(\Psi) \geq 2$ is therefore $1 - o(n^{-1}) - o(1) = 1 - o(1)$. \square

With this, we can show that a linear number of the Ψ_i require multiple responses with probability $1 - o(1)$. This will be enough to give a large lower bound on $\text{cost}(Q(n, m, c))$.

Lemma 7.7. *For each $1 \leq i \leq n$, let ψ_i be a set of cn 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 $Y_i = \{y_i^1, \dots, y_i^n\}$. Further suppose that $m \leq \log_2(n)$ and c, l are any constants with $1 < c < 2$, $l < 1$. With high probability at least ln of the Ψ_i have $\text{cost}(\Psi_i) \geq 2$.*

Proof. In Lemma 7.3, we saw that with high probability $\Psi_i = \exists Y_i \forall X_i. \psi_i$ is false for every $1 \leq i \leq n$, so $\text{cost}(\Psi_i)$ is defined for all Ψ_i . By Lemma 7.6, for each Ψ_i , the probability that $\text{cost}(\Psi_i) \geq 2$ is $1 - o(1)$. Using the Hoeffding bound on the sum of independent Bernoulli random variables, the probability that fewer than ln of the Ψ_i satisfy $\text{cost}(\Psi_i) \geq 2$ is at most

$$\exp\left(-2(1-l-o(1))^2 n\right)$$

which for sufficiently large n can be upper bounded by

$$\exp\left(-2(1-l')^2 n\right)$$

for some constant $l' < 1$. Thus with probability $1 - o(1)$ at least ln of the Ψ_i have cost at least 2. \square

Lemma 7.7 shows that in the randomly generated QBF $Q(n, m, c) \equiv \bigvee_i \Psi_i$, for suitable values of m and c , the Ψ_i are all false and with high probability, a linear proportion of them have $\text{cost}(\Psi_i) \geq 2$. With a slightly more careful choice of m , these two properties suffice to show a cost lower bound for $Q(n, m, c)$. Unfortunately, we cannot obtain $\text{cost}(Q(n, m, c))$ simply by multiplying $\text{cost}(\Psi_i)$ for each $i \in [n]$, as responses on $\text{vars}_{\forall}(\Psi_i)$ may now vary depending on the assignment of variables in some other Ψ_j . Instead, we use the fact that if $\text{cost}(\Psi_i) \geq 2$, then for any response β_i there is some existential assignment for which β_i is *not* a winning response. Using these, for any response β for $Q(n, m, c)$, we construct a large set of existential assignments for which β is not a winning response.

Lemma 7.8. *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 $\text{cost}(Q(n, m, c)) = 2^{\Omega(n^\epsilon)}$.*

Proof. For sets $Y_i = \{y_i^1, \dots, y_i^n\}$, $X_i = \{x_i^1, \dots, x_i^m\}$, with $m \leq (1 - \epsilon) \log_2(n)$,

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

where the clauses C_i^j are chosen uniformly at random to contain two literals on variables in Y_i and a literal on a variable of X_i . For each $1 \leq i \leq n$, define $\Psi_i = \exists Y_i \forall X_i \cdot \bigwedge_{j=1}^{cn} C_i^j$. Let $0 < l < 1$ be a constant. By Lemma 7.7, with probability $1 - o(1)$, all the Ψ_i are false, and at least ln of the Ψ_i have $\text{cost}(\Psi_i) \geq 2$. It therefore suffices to show that if all the Ψ_i are false and $\text{cost}(\Psi_i) \geq 2$ holds for at least ln of the Ψ_i , then $Q(n, m, c)$ is false and $\text{cost}(Q(n, m, c)) \geq 2^{\Omega(n^\epsilon)}$.

If each Ψ_i is false, there is some winning strategy $S_i : \langle Y_i \rangle \rightarrow \langle X_i \rangle$ for each $i \in [n]$. Define $S : \langle Y_1, \dots, Y_n \rangle \rightarrow \langle X_1, \dots, X_n \rangle$ by $S(\alpha_1, \dots, \alpha_n) = (S_1(\alpha_1), \dots, S_n(\alpha_n))$. For any $\alpha \in \langle Y_1, \dots, Y_n \rangle$, restricting $Q(n, m, c)$ by α and $S(\alpha)$ gives

$$Q(n, m, c)[\alpha][S(\alpha)] = \exists t_1 \dots t_n \cdot \bigwedge_{i=1}^n \bigwedge_{j=1}^{cn} \left(\neg t_i \vee C_i^j[\alpha|_{Y_i}][S_i(\alpha|_{Y_i})] \right) \wedge \bigvee_{i=1}^n t_i$$

but by definition of the strategies S_i , $C_i^j[\alpha|_{Y_i}][S_i(\alpha|_{Y_i})] = \perp$ for some $1 \leq j \leq cn$, and so

$$Q(n, m, c)[\alpha][S(\alpha)] = \exists t_1 \dots t_n \cdot \bigwedge_{i=1}^n \neg t_i \wedge \bigvee_{i=1}^n t_i$$

which is clearly unsatisfiable. Since S is a winning \forall -strategy for $Q(n, m, c)$, $Q(n, m, c)$ is false if all the Ψ_i are false for each $i \in [n]$.

It remains to show that $\text{cost}(Q(n, m, c)) \geq 2^{\Omega(n^\epsilon)}$. We may assume that at least ln of the Ψ_i do not have constant winning \forall -strategies. Without loss of generality, we further assume that these are Ψ_1, \dots, Ψ_{ln} , and that all winning \forall -strategies for $Q(n, m, c)$ we consider assign the variables of X_{ln+1}, \dots, X_n according to some constant winning \forall -strategy for $\Psi_{ln+1}, \dots, \Psi_n$. We therefore restrict our attention to strategies which are winning \forall -strategies for Ψ_1, \dots, Ψ_{ln} .

Since $|X_i| \leq (1 - \epsilon) \log_2(n)$, we can list the possible responses in each $\langle X_i \rangle$ as $\langle X_i \rangle = \{\beta_1^i, \dots, \beta_N^i\}$, where $N = 2^m \leq n^{(1-\epsilon)}$.

Let $B = \text{rng}(S)$ for some winning \forall -strategy S for $Q(n, m, c)$. To lower bound $\text{cost}(Q(n, m, c))$, we need to show a lower bound on $|B|$. Given we assume S is constant on $\Psi_{ln+1}, \dots, \Psi_n$, we can consider each $\beta \in B$ as an assignment in $\langle X_1, \dots, X_{ln} \rangle$, i.e. $B \subseteq \{(\beta_{j_1}^1, \dots, \beta_{j_{ln}}^{ln}) : j_1, \dots, j_{ln} \in [N]\}$. As B is the image of a winning \forall -strategy, it contains a winning response β for every assignment $\alpha \in \langle Y_1, \dots, Y_{ln} \rangle$. In this case a winning response for α is some β such that $\Psi_i[\alpha][\beta]$ is false for *every* $1 \leq i \leq ln$.

For each $1 \leq i \leq ln$, Ψ_i does not have a constant winning \forall -strategy. For any $\beta_j^i \in \langle X_i \rangle$, there is some assignment $\alpha_j^i \in \langle Y_i \rangle$ such that β_j^i is not a winning response to α_j^i for Ψ_i . That is, for each β_j^i , there is some α_j^i such that $\Psi_i[\alpha_j^i][\beta_j^i] = \top$, else β_j^i would define a constant winning \forall -strategy for Ψ_i and $\text{cost}(\Psi_i) = 1$. We now use these α_j^i to construct a multiset of existential assignments for which any response β is only a winning response to a small subset.

Define the multiset A , containing elements of $\langle Y_1, \dots, Y_{ln} \rangle$, as

$$A = \left\{ (\alpha_{j_1}^1, \dots, \alpha_{j_{ln}}^{ln}) : (j_1, \dots, j_{ln}) \in [N]^{ln} \right\}$$

Note that α_j^i and $\alpha_{j'}^i$ need not be distinct for $j \neq j'$, so defining A to be a multiset ensures $|A| = N^{ln}$. Given a response $\beta \in B$, we bound the size of the multiset

$$A_\beta = \{ \alpha \in A : \Psi_i[\alpha][\beta] = \perp \text{ for all } 1 \leq i \leq ln \}$$

the set of all assignments in A for which β is a winning response, counted with their multiplicity in A .

For any assignment $\beta \in B$, we know $\beta = (\beta_{j_1}^1, \dots, \beta_{j_{ln}}^{ln})$ for some $j_1, \dots, j_{ln} \in [N]$. If β is a winning response to α , then $\Psi_i[\alpha|_{Y_i}][\beta|_{X_i}] = \perp$ for all $1 \leq i \leq ln$. Since $\beta|_{X_i} = \beta_{j_i}^i$, by the definition of $\alpha_{j_i}^i$, $\alpha|_{Y_i} \neq \alpha_{j_i}^i$ for all $1 \leq i \leq ln$, as $\Psi_i[\alpha_{j_i}^i][\beta_{j_i}^i] = \top$. The restriction $\alpha|_{Y_i} \neq \alpha_{j_i}^i$ restricts the set A_β to

$$A_\beta \subseteq \left\{ (\alpha_{k_1}^1, \dots, \alpha_{k_{ln}}^{ln}) : (k_1, \dots, k_{ln}) \in [N]^{ln}, j_i \neq k_i \text{ for all } 1 \leq i \leq ln \right\}$$

In particular, we see that $|A_\beta| \leq (N-1)^{ln}$.

For any $\alpha \in A$, let $\beta = S(\alpha) \in B$. Since S is a winning \forall -strategy, β is a winning response to α , or equivalently $\alpha \in A_\beta$. By definition, $A_\beta \subseteq A$ for each $\beta \in B$, so it is clear that $A = \bigcup_{\beta \in B} A_\beta$. Comparing the cardinalities of these sets gives $N^{ln} \leq \sum_{\beta \in B} |A_\beta| \leq |B|(N-1)^{ln}$, and so $|B| \geq \left(\frac{N}{N-1}\right)^{ln}$. For $N > 1$, this is a monotonically decreasing function in N , and $N \leq n^{(1-\epsilon)}$, so for sufficiently large n ,

$$|B| \geq \left(\frac{N}{N-1}\right)^{ln} \geq \left(\frac{n^{(1-\epsilon)}}{n^{(1-\epsilon)}-1}\right)^{ln} = \left(1 + \frac{1}{n^{(1-\epsilon)}}\right)^{ln} = \left(\left(1 + \frac{1}{n^{(1-\epsilon)}}\right)^{n^{(1-\epsilon)}}\right)^{ln^\epsilon} = 2^{\Omega(n^\epsilon)}$$

since for large n , $\left(1 + \frac{1}{n}\right)^n \geq 2$. We conclude that $|\text{rng}(S)| \geq 2^{\Omega(n^\epsilon)}$ for any winning \forall -strategy S . There is only one block of universal variables in $Q(n, m, c)$, and so

$$\text{cost}(Q(n, m, c)) = \min\{|\text{rng}(S)| : S \text{ is a winning } \forall\text{-strategy for } Q(n, m, c)\} \geq 2^{\Omega(n^\epsilon)}$$

We have shown that if all the Ψ_i are false, then $Q(n, m, c)$ is false, and further that if at least ln of the Ψ_i have no constant winning \forall -strategy, then $\text{cost}(Q(n, m, c)) \geq 2^{\Omega(n^\epsilon)}$. Lemma 7.7 states that these conditions hold with probability $1 - o(1)$, and this completes the proof. \square

Lemma 7.8 proves that, for the appropriate values of m and c , the QBFs $Q(n, m, c)$ are false and have large cost with probability $1 - o(1)$. It is then a simple application of Size-Cost-Capacity and the capacity upper bounds shown in Section 6 to show lower bounds on $Q(n, m, c)$ with high probability.

Theorem 7.9. *Let $1 < c < 2$ be a constant, and let $m \leq (1 - \epsilon) \log_2(n)$ for some constant $\epsilon > 0$. With high probability, the randomly generated QBF $Q(n, m, c)$ is false, and any QU-Res, CP+ \forall red or PCR+ \forall red refutation of $Q(n, m, c)$ requires size $2^{\Omega(n^\epsilon)}$.*

As previously, the greater capacity of lines in Frege+ \forall red does allow for short proofs of $Q(n, m, c)$ whenever it is false. Refuting any individual false Ψ_i is easy, even for QU-Res. Applying \forall -reduction to each clause results in an unsatisfiable 2-SAT instance, which has a

linear size resolution refutation. This immediately gives a short Frege+ \forall red proof for any false $Q(n, m, c)$, by deriving $\bigvee_{i=1}^n \Psi_i$, and then refuting each Ψ_i in turn.

8. EASY LOWER BOUNDS FOR THE FORMULAS OF KLEINE BÜNING ET AL.

We conclude with a new proof of the lower bounds on the prominent formulas of Kleine Büning et al. [KKF95] using Size-Cost-Capacity.

Definition 8.1 (Kleine Büning et al. [KKF95]). The formulas $\kappa(n)$ are defined as

$$\kappa(n) := \exists y_0 (\exists y_1 \exists y'_1 \forall u_1) \dots (\exists y_k \exists y'_k \forall u_k) \dots (\exists y_n \exists y'_n \forall u_n) (\exists y_{n+1} \dots y_{n+n}) \cdot \bigwedge_{i=1}^{2n} C_i \wedge C'_i$$

where the matrix contains the clauses

$$\begin{aligned} C_0 &= \{\neg y_0\} & C'_0 &= \{y_0, \neg y_1, \neg y'_1\} \\ C_k &= \{y_k, \neg u_k, \neg y_{k+1} \neg y'_{k+1}\} & C'_k &= \{y'_k, u_k, \neg y_{k+1}, \neg y'_{k+1}\} \\ C_n &= \{y_n, \neg u_n, \neg y_{n+1}, \dots, \neg y_{n+n}\} & C'_n &= \{y'_n, u_n, \neg y_{n+1}, \dots, \neg y_{n+n}\} \\ C_{n+t} &= \{\neg u_t, y_{n+t}\} & C'_{n+t} &= \{u_t, y_{n+t}\} \end{aligned}$$

with $1 \leq k \leq n-1$ and $1 \leq t \leq n$.

We also define the QBF $\lambda(n)$ constructed by adding the universal variables v_k for each $1 \leq k \leq n$, quantified in the same block as u_k . The matrix of $\lambda(n)$ contains the clauses D_i, D'_i , where each D_i, D'_i consists of the literals in the corresponding C_i, C'_i , but for each literal on some u_k , we add the matching literal on v_k . This is essentially ‘doubling’ the variables u_k with the matching variables v_k . The effect of this is to prevent any resolution steps being possible on universal variables before the variables can be \forall -reduced.

In [KKF95, BCJ15], it was shown that $\kappa(n)$ requires proofs of size 2^n in Q-Res, which is QU-Res in which universal variables cannot be used as resolution pivots. This lower bound immediately transfers to the same lower bound for $\lambda(n)$ in QU-Res [BWJ14]. As one of the first QBF lower bounds to be shown, these formulas have been the subject of much attention in the study of QBF proof complexity (for examples, see [Egl16, BCJ15, BWJ14, LES16]).

Showing a lower bound for $\kappa(n)$ in Q-Res is equivalent to showing a lower bound for $\lambda(n)$ in QU-Res. It can be assumed in both proof systems that \forall -reductions are performed whenever possible, and so all clauses in the shortest QU-Res proof either contain matching literals on u_k and v_k , or contain no literal on either of them. Any two such clauses cannot be used in a resolution step on a universal variable u_k , as the resulting clause would contain both v_k and $\neg v_k$. All clauses derived from this clause will contain v_k and $\neg v_k$, until a \forall -reduction reduces the clause to \top . The shortest QU-Res proof of $\lambda(n)$ therefore contains no resolution steps on universal pivots, and so the same steps can be used to produce a Q-Res proof of $\kappa(n)$.

We use Size-Cost-Capacity to prove a QU-Res lower bound for an even weaker QBF than $\lambda(n)$, which is obtained by quantifying all the variables v_k in the rightmost universal block. This allows us to give a cost lower bound using this block, which in turn gives the proof size lower bound.

Proposition 8.2. *The QBF $\lambda'(n) := \exists y_0 y_1 y'_1 \forall u_1 \dots \exists y_n y'_n \forall u_n v_1 \dots v_n \exists y_{n+1} \dots y_{n+n} \cdot \bigwedge_{i=1}^{2n} D_i \wedge D'_i$ has cost 2^n .*

Proof. We consider the response of any winning \forall -strategy to the 2^n distinct assignments in the set $A = \{\alpha \in \langle \{y_1, y'_1, \dots, y_n, y'_n\} \rangle : \alpha(y_k) \neq \alpha(y'_k) \text{ for all } 1 \leq k \leq n\}$. Any assignment in A forces a winning \forall -strategy S to respond by setting $u_k = y'_k$. If not, then all clauses C_i, C'_i for $i \leq k$ would be satisfied, and the further assignment $y_j = y'_j = 1$ for all $j > k$ would satisfy the matrix.

It remains to show that responding with $v_k = u_k$ is the only possible response to any $\alpha \in A$ for a winning \forall -strategy. We demonstrate this in the case of the assignment α where $\alpha(y_k) = 1, \alpha(y'_k) = 0$ for all $1 \leq k \leq n$, but other assignments in A are similar. Restricting by the assignment α , as well as β , where $\beta(u_k) = \alpha(y'_k)$ as shown above, the restricted matrix contains the clauses

$$\begin{aligned} D'_n|_{\alpha \cup \beta} &= \{v_n, \neg y_{n+1}, \dots, \neg y_{n+n}\} \\ D'_{n+t}|_{\alpha \cup \beta} &= \{v_t, y_{n+t}\} \quad \text{for each } 1 \leq t \leq n. \end{aligned}$$

If $S_n(\alpha)$ sets any $v_k = 1$, then the matrix is clearly satisfiable by setting $y_{n+k} = 0$, and $y_{n+j} = 1$ for all $j \neq k$. There is therefore a unique response on the v_k for $S_n(\alpha)$, which is to set $v_k = u_k = y'_k$. It is clear that there is a similar unique response for any $\alpha \in A$. We conclude that $|\text{rng}(S_n)| = 2^n$, whence $\text{cost}(\lambda'(n)) = 2^n$. \square

We therefore obtain the following hardness result, which was known for QU-Res [BWJ14], but also lifts to CP+ \forall red and PCR+ \forall red.

Corollary 8.3. *Any QU-Res, CP+ \forall red or PCR+ \forall red proof of $\lambda(n)$ requires size $2^{\Omega(n)}$.*

Proof. The only difference between $\lambda'(n)$ and $\lambda(n)$ is the order in which the variables are quantified. As the variables v_k are quantified further right in $\lambda'(n)$, any refutation of $\lambda(n)$ in any of these proof systems is also a refutation of $\lambda'(n)$. It is therefore sufficient to show a lower bound for refutations of $\lambda'(n)$ in QU-Res, CP+ \forall red or PCR+ \forall red, which is an immediate consequence of Proposition 8.2 and the results of Section 6. \square

This lower bound for QU-Res also yields the lower bound on Q-Res proofs of $\kappa(n)$, previously shown in [KKF95, BCJ15]. As well as providing a relatively simple proof of the hardness of these formulas, our technique also offers some insight as to why these formulas are hard. As the strategy for each variable u_k is simple to compute in even very restricted models of computation, and the proof size lower bounds do not arise from propositional lower bounds, the lower bounds on $\kappa(n)$ and $\lambda(n)$ seemed to be something of an anomaly among QBF proof complexity lower bounds [BHP17]. Here we have shown that the lower bound ultimately arises from the cost of the formulas, although this is slightly obfuscated by some rearrangement of the quantifier prefix.

9. CONCLUSIONS AND OPEN PROBLEMS

By formalising the conditions on P in the construction of $\mathsf{P}+\forall$ red, we developed a new technique for proving QBF lower bounds in $\mathsf{P}+\forall$ red. We demonstrated that the technique is not restricted to a few carefully constructed QBFs, but is in fact applicable to a large class of randomly generated formulas. 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 $P+\forall\text{red}$, and future work will likely establish the hardness of random QBFs in even stronger systems (for example in the expansion based calculus IR-calc [BCJ14]).

Recent work on reasons for hardness in QBF proof systems [BHP17] proposed a new model that precisely characterises the notion of non-genuine hardness in $P+\forall\text{red}$. In that model, the system $\Sigma_1^P\text{-}P+\forall\text{red}$ utilises an P oracle to derive any propositional implicant in a single step, whereby proof size can be measured simply in terms of universal reduction. Hence, by counting universal reduction steps, the Size-Cost-Capacity Theorem in fact gives an absolute proof-size lower bound in $\Sigma_1^P\text{-}P+\forall\text{red}$. We note that our technique also applies to the so-called semantic derivations appearing in [Kra02], since the use of an P oracle guarantees the simulation of any propositional semantic inference.

We also introduced an interesting new formula family, the equality formulas, arguably the simplest known hard QBFs. When considering QBF proof complexity lower bounds, particularly in $P+\forall\text{red}$ 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 essentially as 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.

Future Perspectives. In strong proof systems such as $\text{Frege}+\forall\text{red}$, superpolynomial proof size lower bounds can be completely characterised: they are either a propositional lower bound or a circuit lower bound [BP16]. All QBFs we have considered have no underlying propositional hardness, and winning \forall -strategies can be computed by small circuits, even in very restricted circuit classes. As such, all these QBFs are easy for $\text{Frege}+\forall\text{red}$.

However, for weaker proof systems, such as QU-Res and the others we have discussed, propositional hardness and circuit lower bounds alone are not the complete picture (cf. [BHP17]). In particular, the lower bounds we have shown using Size-Cost-Capacity do not fit into either class. That this 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 the hardness of QBFs in those proof systems where the above dichotomy does not hold. We believe this represents significant progress towards a similar characterisation of lower bounds for these proof systems.

Open Problems. Our work raises several interesting open problems regarding the applicability of our technique to further QBF systems. Firstly, what is the status of Size-Cost-Capacity with respect to long-distance Q-resolution? It is quite simple to construct short proofs of the equality formulas in LD-Q-Res, which demonstrates that the technique as presented in this paper will not lift there, but the underlying ideas could potentially be adapted. The question applies also to expansion-based QBF systems. Finally, the complexity of our random formulas outside of the $P+\forall\text{red}$ paradigm remains an interesting open problem.

ACKNOWLEDGMENTS

Research was supported by a grant from the John Templeton Foundation (grant no. 60842).

REFERENCES

- [ABSRW02] Michael Alekhovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, 2002.
- [AR01] Michael Alekhovich 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.
- [BBC16] 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.
- [BBH18] Olaf Beyersdorff, Joshua Blinkhorn, and Luke Hinde. Size, cost, and capacity: A semantic technique for hard random QBFs. In *Proc. Conference on Innovations in Theoretical Computer Science (ITCS'18)*, 2018.
- [BCJ14] 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.
- [BCJ15] 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.
- [BCMS17] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Feasible interpolation for QBF resolution calculi. *Logical Methods in Computer Science*, 13, 2017.
- [BCMS18a] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Are short proofs narrow? QBF resolution is *not* so simple. *ACM Transactions on Computational Logic*, 19(1):1:1–1:26, 2018.
- [BCMS18b] Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Understanding cutting planes for QBFs. *Information and Computation*, 262:141–161, 2018.
- [BCS17] Olaf Beyersdorff, Leroy Chew, and Karteek Sreenivasaiiah. A game characterisation of tree-like Q-resolution size. *Journal of Computer and System Sciences*, 2017. in press.
- [BGL13] Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A characterization of tree-like resolution size. *Information Processing Letters*, 113(18):666–671, 2013.
- [BHP17] 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.
- [BI10] Eli Ben-Sasson and Russell Impagliazzo. Random CNFs are hard for the polynomial calculus. *Computational Complexity*, 19(4):501–519, 2010.
- [BJ12] Valeriy Balabanov and Jie-Hong R. Jiang. Unified QBF certification and its applications. *Formal Methods in System Design*, 41(1):45–65, 2012.
- [BKS14] 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.
- [BLB10] 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.
- [BM08] 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.
- [BP96] 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.
- [BP16] 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.

- [Bus12] Samuel R. Buss. Towards NP-P via proof complexity and search. *Ann. Pure Appl. Logic*, 163(7):906–917, 2012.
- [BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.
- [BWJ14] 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.
- [CB16] Nadia Creignou and Daniel Le Berre, editors. *International Conference on Theory and Practice of Satisfiability Testing (SAT)*, volume 9710 of *Lecture Notes in Computer Science*. Springer, 2016.
- [CCT87] William Cook, Collette R. Coullard, and György Turán. On the complexity of cutting-plane proofs. *Discrete Applied Mathematics*, 18(1):25–38, 1987.
- [CDER15] 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.
- [CEI96] Matthew Clegg, Jeff Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In Gary L. Miller, editor, *ACM Symposium on Theory of Computing (STOC)*, pages 174–183. ACM, 1996.
- [CFG13] 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.
- [Che17] 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.
- [CI05] 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.
- [CN10] Stephen A. Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity*. Cambridge University Press, Cambridge, 2010.
- [CR79] Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, 1979.
- [CR92] 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.
- [DHK05] 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.
- [DKMP08] 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.
- [dIV98] Wenceslas Fernandez de la Vega. On random 2-SAT (revised version). preprint, 1998.
- [dIV01] Wenceslas Fernandez de la Vega. Random 2-SAT: results and problems. *Theoretical Computer Science*, 265(1-2):131–146, 2001.
- [Egl16] Uwe Egly. On stronger calculi for QBFs. In Creignou and Berre [CB16], pages 419–434.
- [EKLP17] 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.
- [ELW13] Uwe Egly, Florian Lonsing, and Magdalena Widl. Long-distance resolution: Proof generation and strategy extraction in search-based QBF solving. In *International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR)*, pages 291–308, 2013.
- [FLM⁺13] Yuval Filmus, Massimo Lauria, Mladen Miksa, Jakob Nordström, and Marc Vinyals. Towards an understanding of polynomial calculus: New separations and lower bounds - (extended abstract). In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, pages 437–448, 2013.

- [FP83] 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.
- [FPPR17] Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere. Random CNFs are hard for cutting planes. *Computing Research Repository*, abs/1703.02469, 2017.
- [Gel12] 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.
- [GGB11] 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.
- [GMN09] 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.
- [Goe96] Andreas Goerdt. A threshold for unsatisfiability. *Journal of Computer and System Sciences*, 53(3):469–486, 1996.
- [Hak85] Amin Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
- [HP17] Pavel Hrubes and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. *Electronic Colloquium on Computational Complexity*, 24:42, 2017.
- [Jeř05] Emil Jeřábek. *Weak pigeonhole principle, and randomized computation*. PhD thesis, Faculty of Mathematics and Physics, Charles University, Prague, 2005.
- [JM15] Mikoláš Janota and João Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theoretical Computer Science*, 577:25–42, 2015.
- [KKF95] Hans Kleine Büning, Marek Karpinski, and Andreas Flögel. Resolution for quantified Boolean formulas. *Information and Computation*, 117(1):12–18, 1995.
- [KPS⁺09] Roman Kontchakov, Luca Pulina, Ulrike Sattler, Thomas Schneider, Petra Selmer, Frank Wolter, and Michael Zakharyashev. 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.
- [Kra95] Jan Krajčiček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*, volume 60 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, Cambridge, 1995.
- [Kra97] Jan Krajčiček. Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *Journal of Symbolic Logic*, 62(2):457–486, 1997.
- [Kra02] Jan Krajčiček. Interpolation and approximate semantic derivations. *Mathematical Logic Quarterly*, 48(4):602–606, 2002.
- [LES16] Florian Lonsing, Uwe Egly, and Martina Seidl. Q-resolution with generalized axioms. In Creignou and Berre [CB16], pages 435–452.
- [MVC⁺16] 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.
- [MZ09] Sharad Malik and Lintao Zhang. Boolean satisfiability from theoretical hardness to practical success. *Communications of the ACM*, 52(8):76–82, 2009.
- [NPPT09] Massimo Narizzano, Claudia Peschiera, Luca Pulina, and Armando Tacchella. Evaluating and certifying QBFs: A comparison of state-of-the-art tools. *AI Communications*, 22(4):191–210, 2009.
- [PI00] 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.
- [PSS16] Tomáš Peitl, Friedrich Slivovsky, and Stefan Szeider. Long distance Q-resolution with dependency schemes. In Creignou and Berre [CB16], pages 500–518.

- [Pul16] 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.
- [Raz87] Alexander A. Razborov. Lower bounds for the size of circuits of bounded depth with basis $\{\wedge, \oplus\}$. *Mathematical Notes*, 41(4):333–338, 1987.
- [Rin07] Jussi Rintanen. Asymptotically optimal encodings of conformant planning in QBF. In *National Conference on Artificial Intelligence (AAAI)*, pages 1045–1050. AAAI Press, 2007.
- [SB07] 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.
- [Seg07] Nathan Segerlind. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 13(4):417–481, 2007.
- [SM73] 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.
- [Smo87] 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.
- [SS96] 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.
- [Var14] Moshe Y. Vardi. Boolean satisfiability: Theory and engineering. *ACM*, 57(3):5, 2014.
- [Vol99] Heribert Vollmer. *Introduction to Circuit Complexity - A Uniform Approach*. Texts in Theoretical Computer Science. Springer, 1999.