Electronic Colloquium on Computational Complexity, Report No. 137 (2017)

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

Olaf Beyersdorff, Joshua Blinkhorn, and Luke Hinde

School of Computing, University of Leeds, UK

Abstract. As a natural extension of the SAT problem, different proof systems for quantified Boolean formulas (QBF) have been proposed. Many of these extend a propositional system to handle universal quantifiers.

By formalising the construction of the QBF proof system $\mathsf{P}{+}\forall\mathsf{red}$ from a propositional proof system P (Beyersdorff, Bonacina & Chew, ITCS '16), we present a new technique for proving proof size lower bounds in these systems. This lower bound technique relies only on two semantic properties: the *cost* of a QBF, and the *capacity* of a proof. By examining the capacity of proofs in several proof systems, we are able to use this technique to obtain lower bounds in these systems based on cost alone.

As applications of this 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.

1 Introduction

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 [16,23] and bounded arithmetic [22,43].

Proof complexity is intrinsically linked to recent noteworthy innovations in solving, owing to the fact that any decision procedure implicitly defines a *proof system* for the underlying language. Relating the two fields in this way is illuminating for the practitioner; proof-size and proof-space lower bounds correspond directly to best-case running time and memory consumption for the corresponding solver. Indeed, proof complexity theory has become the main driver for the asymptotic comparison of practical 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 exceptional strength of modern implementations. Cutting-edge advances in solving, spearheaded by unparalleled progress in Boolean satisfiability (SAT), appear to provide a means for the efficient solution of computationally hard problems [56]. Contemporary SAT solvers routinely dispatch instances in millions of clauses [45], and are effectively employed as **NP**-oracles in more complex settings [46]. 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) [52]. Besides furthering the intense study of resolution and its fragments [16], 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**-

ISSN 1433-8092

complete language [55]. QBF extends propositional logic with existential and universal quantification, and consequently offers succinct encodings of concrete problems from conformant planning [17, 29, 51], ontological reasoning [42], and formal verification [7], amongst other areas [14, 26, 54]. 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 [4, 10, 39].

The semantics of QBF has a neat interpretation as a two-player *evaluation game*. Given a QBF $\Phi = \mathcal{Q} \cdot \phi$, the $\exists$- and $\forall$-players take turns to assign the existential and universal variables of the formula following the order of the quantifier prefix $\mathcal{Q}$. When all variables are assigned, the $\exists$-player wins if the propositional formula ϕ is satisfied; otherwise, the $\forall$-player takes the win. A folklore result states that a QBF is false if and only if the $\forall$-player can win the evaluation game by force; that is, if and only if there exists a winning strategy for the universal player. The concept of *strategy extraction* originates from QBF solving [36], 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 – in real-world applications, certificates provide further useful information [54].

A major paradigm in QBF practice is *quantified conflict-driven clause learning* (QCDCL) [34], a natural extension of CDCL. The vast majority of QBF solvers build upon existing SAT techniques in a similar fashion. Such a notion can hardly be surprising when one considers that an existentially quantified QBF is merely a propositional formula; hence every QBF implementation contains an embedded SAT solver *by default*. 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* ($\mathsf{QU\text{-}Res}$) [33,41], a well-studied QBF proof system that underpins 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, $\mathsf{QU\text{-}Res}$ is identical to resolution, and proof-size lower bounds for the latter lift immediately to the former. From the viewpoint of quantified logic, lower bounds obtained in this way are rightly considered *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 $\mathsf{P}+\forall\mathsf{red}$ [8], which augments the propositional rules of P with the universal reduction rule $\forall\mathsf{red}$. As a consequence, the phenomenon of genuineness extends well beyond resolution. In this paper, in addition to resolution we consider three stronger systems: Cutting Planes (CP), a well-studied calculus that works with linear inequalities; the algebraic system Polynomial Calculus (with Resolution, PCR); and Frege's eponymous 'textbook' system for propositional logic. Their simulation order is depicted in Figure 1.

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

When is a lower bound genuine? Naturally, the aforementioned objections to non-genuine QBF lower bounds may be raised in the abstract setting of $\mathsf{P}+\forall\mathsf{red}$, as that system encompasses the propositional proof system P. Indeed, given any unsatisfiable propositional formulas

[1] The calculus $\mathsf{QU\text{-}Res}$, proposed by Van Gelder in [33], generalises $\mathsf{Q\text{-}Res}$, introduced by Kleine Büning et al. in [41], by allowing resolution over universally quantified pivots.

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 $\mathsf{P}+\forall\mathsf{red}$, but whose hardness stems from the original propositional formulas. That such lower bounds ought to be identified as non-genuine was highlighted in [18] (cf. also [12]).

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

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

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

In propositional proof complexity, random 3-SAT instances, the most commonly studied random CNFs, are relatively well understood. There is a constant r such that if a random CNF on n variables contains more than rn clauses, then the CNF is unsatisfiable with probability approaching 1 [32]; the upper bound for r has regularly been improved (see [27], 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 [5]. Hardness results for random CNFs are also known for Polynomial Calculus [2, 6] and for Cutting Planes [31, 38].

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 [19], for which bounds on the threshold number of clauses needed for a false QBF were shown in [25]. 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.

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 $\mathsf{P}+\forall\mathsf{red}$, 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 will henceforth occupy a prominent place in QBF proof complexity. As our principal application we prove exponential lower-bounds in three concrete $\mathsf{P}+\forall\mathsf{red}$ 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

[2] This notion can be made formal, as in the oracle model of [12].

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

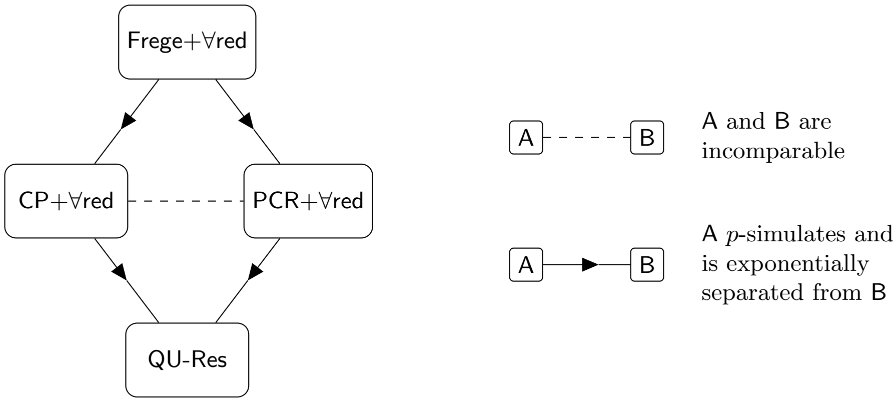

Fig. 1. The simulation order of the four QBF proof systems featured in this paper. A proof system A p-simulates the system B if each B-proof of a formula Φ can be translated in polynomial time into an A-proof of Φ [23]. If neither A nor B p-simulates the other, then they are incomparable.

our technique can be applied to give a simple proof of lower bounds for a family of well known QBFs from [41].

In addition, we also determine exact conditions on a so-called *base system* P by which P+∀red is properly defined and receptive to our method. We detail our contributions below, beginning with the lower bound technique, followed by several applications.

2.1 A new technique for genuine QBF lower bounds

Using an established approach (e.g. [48]), the soundness of P+∀red is proved by demonstrating that a winning strategy for the ∀-player can be extracted from a refutation (Lemma 15). However, with careful construction and analysis of the strategy extraction algorithm, we are able to obtain a much more valuable result.

Given a P+∀red refutation π of a QBF Φ, strategy extraction works by first restricting π according to the ∃-player's move, then collecting the response for the ∀-player from some line in π, and iterating until the evaluation game concludes. We therefore reason as follows: A lower bound on the total number of responses contributed by π, coupled with an upper bound on the number of responses contributed per line, yields a lower bound on the number of lines in the refutation.

In light of this observation, we define the two measures called *cost* and *capacity*. 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 23 (Size-Cost-Capacity Theorem). *Let* P *be a base system, and let* π *be a* P+∀red *refutation of a QBF* Φ*. Then*

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

We also show explicitly that Size-Cost-Capacity works by counting universal reduction steps (Lemma 22), which illustrates that all results obtained by application of our technique are genuine QBF lower bounds in the aforementioned sense.

Moving on, we prove that all QU-Res and CP+∀red refutations have capacity equal to 1 (Propositions 21 and 26). Hence, the Size-Cost-Capacity Theorem tells us that *cost alone* gives an absolute lower bound on proof size there. The case for the QBF version of Polynomial Calculus with Resolution (PCR+∀red) is much more challenging, and requires some linear algebra, owing to the underlying algebraic composition of Polynomial Calculus (see Subsection 6.2). Interestingly, it turns out that the capacity of a refutation there is no greater than its size (Proposition 29), thus proof size is at least the square root of cost. Hence, we obtain the following absolute proof-size lower bounds.

Corollaries 24 and 27. *Let π be a* QU-Res *or* CP+∀red *refutation of a QBF Φ. Then* $|\pi| \geq \mathrm{cost}(\Phi)$.

Corollary 30. *Let π be a* PCR+∀red *refutation of a QBF Φ. Then* $|\pi| \geq \sqrt{\mathrm{cost}(\Phi)}$.

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

2.2 Applications of the technique

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

A. The equality formulas: a non-trivial special case. In order to illustrate by example the concept behind our technique – and before proving the general case – we first show a new exponential lower bound for QU-Res. Our argument rests on two folklore propositions – the component parts of the strategy extraction algorithm – and serves to highlight the important points ahead of the more difficult general case. The result is likely of independent interest, owing to the introduction of an interesting new family of hard QBFs.

Definition 2 (equality formulas). *For $n \in \mathbb{N}$, the n^{th}* equality formula *is*

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

The equality formulas are so called because the only winning strategy for the ∀-player in the evaluation game is as follows: play $u_i = x_i$ for each $i \in [n]$. Consequently the winning strategy is not only unique, it contains all 2^n assignments to the universal variables. These two properties in tandem are largely responsible for the apparent ease with which we prove the following result.

Theorem 5. *Let π be a* QU-Res *refutation of $\Theta(n)$. Then $|\pi| \geq 2^n$.*

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

Closer inspection reveals that this lower bound is of a very specific type – it is a genuine QBF lower bound (the formulas are not harbouring propositional hardness) that does not derive from a circuit lower bound (the winning strategy is not hard to compute in an associated circuit class). In existing QBF literature, the only other example of such a family comes from the famous formulas of Kleine Büning et al. [41] (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.

Once we have shown the technique in its full generality, we then use the exponential cost of the equality formulas (Proposition 18) to conclude that they not only require exponential-size refutations in QU-Res (as shown directly in Theorem 5), but even in the stronger systems CP+∀red and PCR+∀red. They do however have linear-size refutations in Frege+∀red (Proposition 31). It is interesting to note that the short refutations we give have exponential capacity.

B. The first hard random QBFs. For the major application of our technique, we define a class of random QBFs and prove that, with high probability, they are hard in all three systems QU-Res, CP+∀red and PCR+∀red. We generate instances that combine the overall structure of the equality formulas with the literature's existing model of random QBFs from [19].

Definition 32. *For each $1 \leq i \leq n$, let $C_i^1, \dots, C_i^{cn}$ be distinct clauses picked uniformly at random from the set of clauses containing 1 literal from the set $X_i = \{x_i^1, \dots, x_i^m\}$ and 2 literals from $Y_i = \{y_i^1, \dots, y_i^n\}$. Define the randomly generated QBF $Q(n, m, c)$ as:*

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

The specification of how many existential and universal variables each clause should contain is a common and necessary restriction on random QBFs [19, 25]. This prevents the occurrence of a clause containing only universal variables – if such a clause exists, there is a constant size refutation of this clause alone in any P+∀red system. The motivation behind the additional structure in the construction of $Q(n, m, c)$ is that its truth value is equivalent to the disjunction of its 'component parts'; that is $Q(n, m, c) \equiv \bigvee_{i=1}^{n} \Psi_i$, where $\Psi_i := \exists Y_i \forall X_i \cdot \bigwedge_{j=1}^{cn} C_i^j$ for each $i \in [n]$.

These Ψ_i are some of the simplest QBFs one can generate, so $Q(n, m, c)$ is a natural choice of random QBFs. Indeed, the model used to generate the clauses of Ψ_i is also used to generate random QBFs for the evaluation of QBF solvers [15, 49].

Drawing on the existing literature [20, 25, 58], we show that suitable choices of the parameters m and c force each Ψ_i to be false with high probability. The individual Ψ_i are essentially equivalent to a random 2-SAT problem, and this step is just an application of results on the satisfiability of such instances.

Moreover, we also prove a cost lower bound. Perhaps surprisingly, this cost lower bound is constructed by applying results on the unsatisfiability of random 2-SAT instances [58] and the truth of random (1,2)-QCNFs [25]. These results both concern only the truth value of the corresponding formulas, and taken individually seem unrelated to cost. However, by carefully choosing the number of clauses so as to allow the application of both results, we can construct a cost lower bound using the following argument.

The Ψ_i are false with high probability, but rearranging the quantifiers to $\forall X_i \exists Y_i \cdot \bigwedge_{j=1}^{cn} C_i^j$ gives a QBF which is true with probability $1 - o(1)$. In other words, with high probability, the universal response in Ψ_i must depend on the existential assignment. In particular, it must change depending on the existential assignment, and so with probability $1 - o(1)$, linearly many of the Ψ_i require at least two distinct responses in any winning strategy (Lemma 34).

By refining our choice of m slightly, this allows us to conclude that $Q(n,m,c)$, with high probability, is a false QBF with large cost.

Proposition 39. *Let $1 < c < 2$ be a constant, and let $m \leq (1-\epsilon)\log_2(n)$ for some constant $\epsilon > 0$. With probability $1 - o(1)$, $Q(n,m,c)$ is false and $\mathrm{cost}(Q(n,m,c)) = 2^{\Omega(n^{\epsilon})}$.*

Invoking Theorem 23 yields immediate hardness results. The following theorem constitutes the first proof-size lower bounds for random generated formulas in the QBF proof complexity literature. We emphasize that these are genuine QBF lower bounds in the aforementioned sense; they are not merely hard random CNFs lifted to QBF. As for any application of Size-Cost-Capacity, the refutations are large precisely because they require many universal reduction steps.

Theorem 40. *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+∀red *or* PCR+∀red *refutation of $Q(n,m,c)$ requires size $2^{\Omega(n^{\epsilon})}$.*

C. New proofs of known lower bounds. We conclude by using Size-Cost-Capacity to provide a new proof of the hardness of the prominent QBFs of Kleine Büning, Karpinski and Flögel [41]. We consider a common modification of the formulas, denoted by $\lambda(n)$, which consists of 'doubling' each universal variable. This modification is known to lift lower bounds in Q-Res to lower bounds in QU-Res [4], where we can apply Size-Cost-Capacity.

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

Corollary 43. *Any* QU-Res*,* CP+∀red *or* PCR+∀red *proof of $\lambda(n)$ requires size $2^{\Omega(n)}$.*

As QU-Res lower bounds on these modified formulas are shown to be equivalent to Q-Res lower bounds on the original formulas, our technique even proves the original lower bounds from [41] (cf. also [10]), and provides some insight as to the source of this lower bound.

Generality of applicability. In order to present Size-Cost-Capacity in total generality, we take the concept of P+∀red (introduced in [8] for a hierarchy of Frege systems) and formalise exact conditions on P yielding a sound and complete QBF proof system. We identify three natural properties that are sufficient: (a) The derivable axioms are 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 7). On account of the low-level generality, the following theorem requires a non-trivial proof.

Theorem 11. *If* P *is a base system, then* P+∀red *is a sound and complete QBF proof system.*

Formalising the framework of base systems thus renders Size-Cost-Capacity applicable to the complete spectrum of P+∀red systems. All the concrete propositional calculi considered in this work (i.e. those appearing in Figure 1) are demonstrably base systems.

2.3 Relation to previous work

Strategy extraction for QBF lower bounds has been explored previously by exploiting connections to circuit complexity [8, 10, 13]. In particular, [8] established tight relations between circuit and proof complexity, lifting even strong circuit lower bounds for $\mathbf{AC^0}[p]$ circuits [50,53] to QBF lower bounds for $\mathbf{AC^0}[p]$-Frege+$\forall$red [8], which is unparalleled in the propositional domain. In fact, for strong proof systems such as Frege+$\forall$red, this strategy extraction technique is sufficient to prove any genuine QBF lower bound, in the sense that any superpolynomial lower bound for Frege+$\forall$red arises either due to a lower bound for Frege, or due to a lower bound for Boolean circuits [13]. 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 [12]. 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 [8, 10, 13], of course, is the rarity of superpolynomial lower bounds from circuit complexity [59], 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.

2.4 Innovations and future perspectives

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

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

Size-Cost-Capacity also opens new research avenues concerning the reasons for QBF hardness – a topic that is currently insufficiently understood. By presenting the technique for P+$\forall$red in general, we are able (in certain cases) to associate large proofs with low capacity and short proofs with high capacity. This goes some way towards explaining why the observed dichotomy in Frege+$\forall$red [13] breaks down for weaker systems such as QU-Res: the reduced capacity *introduces a new form of genuine QBF hardness*. As such, our work opens the door for a better understanding, and makes steps towards the complete characterisations of reasons for hardness that are currently lacking in the literature.

2.5 Organisation of the paper

After dealing with preliminaries in Section 3, we give the direct proof of hardness for the equality formulas in Section 4. This is a relatively straightforward example of our new lower bound technique, and is presented first to help illustrate the more general arguments of Section 5.

In Section 5, we give all the details of our lower bound technique and prove the Size-Cost-Capacity Theorem. Subsections 5.1 and 5.2 provide low-level details pertaining to the definition and completeness of P+$\forall$red. If the reader wishes to skip such details, and go straight to the lower-bound technique, they might begin with strategy extraction in Subsection 5.3.

Upper bounds on capacity for CP+$\forall$red and PCR+$\forall$red are the subject of Section 6, and the material on random QBFs appears in Section 7. In Section 8, we provide a new proof of some well-known QBF lower bounds using our new technique, and finally, we offer some concluding thoughts in Section 9.

3 Preliminaries

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

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

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

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

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

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

Proposition 1 (folklore). *A QBF is false if and only if it has a winning $\forall$-strategy.*

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

---

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

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

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

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

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

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.

QU-Resolution (QU-Res) [33, 41] is a resolution-based proof system for QBFs of the form $\Phi = \mathcal{Q} \cdot \phi$, where ϕ is a CNF. The calculus supplements resolution with a *universal reduction rule* which allows (literals in) universal variables to be removed from a clause C provided that they are right of all existentials in C with respect to $\mathcal{Q}$. Tautological clauses are explicitly forbidden; for any variable x, one may not derive a clause containing both x and $\neg x$. The rules of QU-Res are given in Figure 2. Note that we choose to include weakening of clauses as a valid inference rule. Whereas this is not conventional,[7] it is justified, since its inclusion is consistent with the overall narrative of this paper.

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

4 A QU-Res lower bound for the equality formulas

In this section, we introduce the *equality formulas* and sketch 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. For that reason, technical details and formal proofs are omitted; formal proofs of the arguments used here are encompassed by those in the following section on strategy extraction in P+∀red.

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

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

[7] A weakening step does not occur in the shortest QU-Res refutation of any QBF.

Note that the propositional part of $\Theta(n)$ is the CNF consisting of the *long clause* $\{t_1, \ldots, t_n\}$ 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 $\{t_1, \ldots, t_n\}$. 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 of all this 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 $\Theta(n)$ is exactly the set $\langle\{u_1, \ldots, u_n\}\rangle$. Now, given a refutation π of $\Theta(n)$, we will prove by appeal to strategy extraction that each element of that range appears in π as a subset of some clause. Since no two elements can be subsets of the same clause (that would produce a universal tautology), the size of π is at least $|\langle\{u_1, \ldots, u_n\}\rangle| = 2^n$.

Overview of round-based strategy extraction. Strategy extraction is an important QBF paradigm that was motivated by solving certification (cf. [36, 47]), and subsequently received much attention in the literature [3, 8, 30, 48]. In this paper, we follow the algorithm given in [36], which for the sake of clarity we refer to as *round-based strategy extraction*.

Given a QU-Res refutation of a PCNF, round-based strategy extraction is an iterative procedure that returns a winning $\forall$-strategy. During the course of the game, the $\forall$-player maintains a restriction of the refutation, from which her winning moves may be determined. She need only do the following two things:

(a) restrict the refutation by the $\exists$-player's move.
(b) 'read off' the response from the restricted refutation.

These two steps are simply repeated round by round until the game concludes.[8]

The correctness of the procedure rests on two corresponding propositions, which originate from [36]:[9] (a) QU-Res refutations are preserved by existential restrictions (Proposition 3); (b) a winning response can be determined algorithmically from the refutation (Proposition 4). To prove the hardness of $\Theta(n)$, we do not need to formalise QU-Res strategy extraction and prove its correctness; it suffices to argue directly from these two propositions.

Direct proof of hardness. We briefly describe how to restrict a refutation by an existential assignment. For each QU-Res refutation π of a PCNF Φ, and each partial assignment α to the existential variables of Φ, we define the restricted refutation $\pi[\alpha]$ as follows: Let C_i be the first clause in π for which $C_i[\alpha] = \bot$, and let π_i be the subderivation of C_i. Then $\pi[\alpha]$ is the sequence obtained by restricting each clause in π_i by α, while removing all satisfied clauses. QU-Res refutations are preserved by existential restrictions in the following sense.

Proposition 3. *Let π be a* QU-Res *refutation of a PCNF Φ, and let α be a partial assignment to the existential variables of Φ. Then $\pi[\alpha]$ is a* QU-Res *refutation of $\Phi[\alpha]$.*

Now, given a QBF whose first block U is universal, a winning move for the universal player on block U can be determined easily from a QU-Res refutation, as follows. If the final step is a universal reduction, select an assignment that falsifies all the reduced literals (this must be

[8] In practice, the refutation is commonly restricted by the $\forall$-player's move as well (cf. [36]), but this is not strictly necessary.

[9] That paper presents analogous results for the weaker system Q-Res.

possible since universal tautologies are disallowed). Otherwise, if the final step is a resolution, any assignment to U is a winning move.

Proposition 4. *Let π be a non-trivial* QU-Res *refutation of a PCNF Φ whose first block U is universal, and let β be a total assignment to U. If β falsifies the universal literals in the penultimate clause of π, then $\Phi[\beta]$ is false.*

To prove the lower bound we claim the following: In a refutation π of $\Theta(n)$, each total assignment β to the universal variables appears as a subset of some clause. Since tautologies are forbidden, the claim implies that π contains at least 2^n clauses.

The claim is established from our two propositions and properties of $\Theta(n)$ as follows: Let U be the universal variables of $\Theta(n)$, and let β be any total assignment to U. Now consider the unique total assignment α to the existential variables $\{x_1, \dots, x_n\}$ that satisfies $\alpha(x_i) \neq \beta(u_i)$ for each $i \in [n]$. By Proposition 3, $\pi[\alpha]$ is a refutation of $\Theta(n)[\alpha]$. It is easy to verify that $\pi(n)$ is non trivial, so let C denote its penultimate clause. Now, by the uniqueness of the countermodel for $\Theta(n)$, the only total assignment β' to U for which $\Theta(n)[\alpha][\beta']$ is false must satisfy $\beta'(u_i) = \alpha(x_i)$; that is, $\beta'(u_i) \neq \beta(u_i)$ for each $i \in [n]$. Moreover, by Proposition 4, β' is the only total assignment to U that falsifies the universal literals in C. It follows that β, represented as a set of literals, is contained in C. The claim follows, since C is contained in a clause of the original proof π by the definition of restriction.

Theorem 5. *Let π be a* QU-Res *refutation of $\Theta(n)$. Then $|\pi| \geq 2^n$.*

In a nutshell, Theorem 5 was proved by equating the minimum refutation size with the cardinality of the range of a winning $\forall$-strategy for $\Theta(n)$. Our argument here was aided by two facts: $\Theta(n)$ has a unique winning $\forall$-strategy and contains a single universal block. Of course, neither fact holds for QBFs in general. Nonetheless, in the following section, we generalise the method to prove an absolute proof-size lower bound for any instance in P+$\forall$red.

5 A new lower bound technique for P+$\forall$red

In this section, we develop a general technique for proof-size lower bounds in P+$\forall$red, by extrapolation from the method of Section 4. In Subsection 5.1, we first describe precisely what we mean by a *line-based propositional proof system* P, and proceed to identify three natural conditions by which P+$\forall$red is a QBF proof system to which our technique applies. The formal definition of P+$\forall$red (adapted from [8]) and a proof of its completeness is provided in Subsection 5.2. In Subsection 5.3, we formally define round-based strategy extraction for P+$\forall$red and prove its correctness. Finally, in Subsection 5.4 we state and prove our central result, the *Size-Cost-Capacity Theorem.*

5.1 Line-based propositional proof systems

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

(a) A set of *lines* $\mathcal{L}_\mathsf{P}$, containing at least the two lines $\top$ and $\bot$ that represent trivial truth and trivial falsity, respectively.
(b) A set of *inference rules* $\mathcal{I}_\mathsf{P}$ and an *axiom function* that maps each propositional formula ϕ to a set of axioms $\mathcal{A}_\mathsf{P}(\phi) \subseteq \mathcal{L}_\mathsf{P}$.

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

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

(c) A *variables function* that maps each line $L \in \mathcal{L}_{\mathsf{P}}$ to a finite set of Boolean variables $\text{vars}(L)$, satisfying $\text{vars}(\top) = \text{vars}(\bot) = \emptyset$. Additionally, $\text{vars}(L) \subseteq \text{vars}(\phi)$ for each line L in a P-derivation from ϕ.[10]

(d) A *restriction operator* (denoted by square brackets) that takes each line $L \in \mathcal{L}_{\mathsf{P}}$, under restriction by any partial assignment τ to $\text{vars}(L)$, to a line $L[\tau] \in \mathcal{L}_{\mathsf{P}}$. If τ is a total assignment, then $L[\tau]$ is either $\top$ or $\bot$. Restriction of L by an arbitrary Boolean assignment σ is defined as the restriction of L by the projection of σ to $\text{vars}(L)$.

The purpose of the restriction operator is to encompass the natural semantics of P. For that reason, we make the natural stipulation that restriction by a total assignment to the variables of a line yields either trivial truth or trivial falsity. We may therefore associate with any line $L \in \mathcal{L}_{\mathsf{P}}$ the Boolean function on $\text{vars}(L)$ that computes the propositional models of L, with respect to the semantics of the restriction operator for P.

Definition 6 (associated Boolean function). *Let* P *be a line-based propositional proof system and let* $L \in \mathcal{L}_{\mathsf{P}}$*. The* associated Boolean function *for* L *is* $B_L : \langle \text{vars}(L) \rangle \to \{0,1\}$*, defined by*

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

Beyond the established notion of *line-based*, we identify three natural properties by which P can be augmented with $\forall$-reduction, yielding a bona fide QBF proof system $\mathsf{P}+\forall\mathsf{red}$. 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.*[11] 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 7 (base system). *A* base system P *is a line-based propositional proof system satisfying the following three properties:*

(a) **Axiomatic equivalence**. *For each propositional formula* ϕ *and each* $\tau \in \langle \text{vars}(\phi) \rangle$*,* $\phi[\tau] = \top$ *iff each* $A \in \mathcal{A}_{\mathsf{P}}(\phi)$ *satisfies* $A[\tau] = \top$*.*

(b) **Inferential equivalence**. *For each set of lines* $\mathcal{L} \subseteq \mathcal{L}_{\mathsf{P}}$ *and each line* $L \in \mathcal{L}_{\mathsf{P}}$*,* L *can be derived from* $\mathcal{L}$ *iff* $\mathcal{L}$ *semantically entails* L*.*

(c) **Restrictive closure**. *For each* $L \in \mathcal{L}_{\mathsf{P}}$ *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) is a base system. The axiomatic equivalence is trivial, as is inferential equivalence, which follows directly from implicational completeness and logical correctness. Taking the conventional definitions of the variable

[10] Note that this does not exclude extended Frege systems (EF), whose lines can be represented as Boolean circuits as in [40, p. 71].

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

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

– β is a partial assignment to the universal variables of Φ.

– each universal in $\text{vars}(\beta)$ is right of each existential in $\text{vars}(L)$, with respect to $\mathcal{Q}$.

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

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.

We conclude the subsection with two useful propositions. From the definition of base system, it does *not* follow that the order of successive restrictions of a line may be ignored (that is, $L[\tau_1][\tau_2]$ does not equal $L[\tau_1 \cup \tau_2]$ in general). However, it is a straightforward consequence of restrictive closure that the associated Boolean function is preserved.

Proposition 8. *Let* P *be a base system, let* $L \in \mathcal{L}_{\mathsf{P}}$*, and let* $\tau_1, \dots, \tau_n$ *be pairwise variable-disjoint, partial assignments to* $\text{vars}(L)$*. Then the associated Boolean functions for* $L[\tau_1]\cdots[\tau_n]$ *and* $L[\bigcup_{i=1}^n \tau_i]$ *are identical.*

Proof. We have $B_{L[\tau_1]\cdots[\tau_n]} = (\cdots(B_L|_{\tau_1})\cdots)|_{\tau_n} = B_L|_{\bigcup_{i=1}^n \tau_i} = B_{L[\bigcup_{i=1}^n \tau_i]}$, by the restrictive closure of P. $\square$

Finally, we show that a base system has the power to express disjunctions; that is, for each assignment to Boolean variables, there exists a line in $\mathcal{L}_{\mathsf{P}}$ which is falsified only by that assignment, or an extension of it.

Proposition 9. *Let* P *be a base system, and let* τ *be an assignment to Boolean variables. Then there exists a line* $L \in \mathcal{L}_{\mathsf{P}}$ *with* $\text{vars}(L) = \text{vars}(\tau)$ *for which* B_L *is zero only at* τ*.*

Proof. Let $\sigma \in \langle \text{vars}(\tau) \rangle$, let $\phi = \bigvee_{l \in \tau} \neg l$ and observe that $\phi[\sigma] = \bot$ iff $\sigma = \tau$. By the axiomatic equivalence of P, there is some $A \in \mathcal{A}_{\mathsf{P}}(\phi)$ for which $A[\sigma] = \bot$ iff $\sigma = \tau$. Since A may be introduced in a P-derivation from ϕ, we have $\text{vars}(A) \subseteq \text{vars}(\tau)$. Aiming for contradiction, suppose that $\text{vars}(A) \subset \text{vars}(\tau)$, and let σ' be the projection of σ to $\text{vars}(A)$. Then, restriction of A by any extension of σ' returns $\bot$, but at least one such extension is in $\langle \text{vars}(\tau) \rangle$ and not equal to τ, a contradiction. It follows that $\text{vars}(A) = \text{vars}(\tau)$, and B_A is zero only at τ. $\square$

5.2 P+∀red: definition and completeness

Universal reduction is a widely used rule of inference in QBF proof systems, by which universal variables may be assigned under certain conditions. More precisely, a line L may be restricted by an assignment to a universal variable u provided it is right of all the existentials in $\text{vars}(L)$, with respect to the prefix of the input QBF. By restrictive closure, the restriction of a line by an assignment to u results in the exclusion of u from the domain of the associated Boolean function. Universal reduction should therefore be viewed as a sound method for deleting universal variables. We state the rule formally in Figure 3; note that we allow multiple reductions in a single step (that is, restriction by a partial assignment) provided that each individual universal variable is eligible.

The primary purpose of universal reduction is to lift a propositional proof system P to a QBF system P+∀red, as in the following definition.

Definition 10 (P+∀red [8]). *Let* P *be a line-based propositional proof system. Then* P+∀red *is the system consisting of the inference rules of* P *in addition to universal reduction, in which references to the input formula ϕ in the rules of* P *are interpreted as references to the propositional part of the input QBF $\mathcal{Q} \cdot \phi$.*

We extend our notation from P to P+∀red in the natural way, denoting the lines available in P+∀red (syntactically equivalent to the lines available in P) by $\mathcal{L}_{\mathsf{P+\forall red}}$, and writing $\mathrm{vars}_\exists(L)$ and $\mathrm{vars}_\forall(L)$ for the subsets of $\mathrm{vars}(L)$ consisting of the existentially and universally quantified variables, respectively. Also, we observe that Res + ∀red and QU-Res are (virtually) identical proof systems,[12] and we will henceforth use the terms interchangeably.

In this subsection and the next, we prove that P+∀red is sound and complete for QBF if P possesses the three properties of a base system.[13]

Theorem 11. *If* P *is a base system, then* P+∀red *is a sound and complete QBF proof system.*

The argument for completeness is straightforward. A winning ∀-strategy S for a false QBF Φ can be represented equivalently as the set $\{\alpha \cup S(\alpha) : \alpha \in \langle \mathrm{vars}_\exists(\Phi)\rangle\}$ of total assignments to the variables of Φ. By Proposition 9, for each such element $\tau = \alpha \cup S(\alpha)$, we can derive from Φ some line in $\mathcal{L}_{\mathsf{P+\forall red}}$ whose associated Boolean function is zero only at τ. By deriving such a line for each $\alpha \in \mathrm{vars}_\exists(\Phi)$, we can faithfully represent the strategy S within a P+∀red-derivation from Φ. Starting at the rightmost block, we successively 'truncate' this representation of the strategy, using ∀-reduction to delete universal variables, and the implicational completeness of P to remove existentials. At the final step, the strategy collapses completely, and we derive trivial falsity.

Lemma 12. *If* P *is a base system, then* P+∀red *is complete.*

Proof. Let S be a winning ∀-strategy for a false QBF $\Phi = \forall U_1 \exists E_1 \cdots \forall U_n \exists E_n \cdot \phi$. Further, let $S_0 = \{\emptyset\}$, and, for each $i \in [n]$, let S_i be the set consisting of the projection of the assignments $\{\alpha \cup S(\alpha) : \alpha \in \langle \mathrm{vars}_\exists(\Phi)\rangle\}$ to the variables $\bigcup_{j=1}^{i}(U_j \cup E_j)$ (informally, S_i is the 'truncation' of S to the first i rounds). Finally, for each assignment τ to Boolean variables, we denote by L_τ some fixed line in $\mathcal{L}_\mathsf{P}$ with $\mathrm{vars}(L_\tau) = \mathrm{vars}(\tau)$ for which B_{L_τ} is zero only at τ. Such a line L_τ exists by Proposition 9.

We will prove by backwards induction on $i \in \{0, \dots, n\}$ that $\Phi \vdash_{\mathsf{P+\forall red}} L_\tau$ for each $\tau \in S_i$. Hence we prove the theorem at the final step $i = 0$, since $L_\emptyset = \bot$.

For the base case $i = n$, we observe that $\phi[\tau] = \bot$ for each $\tau \in S_n$, by definition of winning ∀-strategy. Hence, by the axiomatic equivalence of P, in a P-derivation from ϕ we can introduce, for each $\tau \in S$, some axiom $A_\tau \in \mathcal{A}_\mathsf{P}(\phi)$ for which $A_\tau[\tau] = \bot$. Now, since $A_\tau \vDash L_\tau$, we have $A_\tau \vdash_\mathsf{P} L_\tau$ by the implicational completeness of P. It follows that $\phi \vdash_\mathsf{P} L_\tau$ (and hence $\Phi \vdash_{\mathsf{P+\forall red}} L_\tau$) for each $\tau \in S_n$.

For the inductive step, let $i \in [n]$ and suppose that $\Phi \vdash_{\mathsf{P+\forall red}} L_\tau$ for each $\tau \in S_i$. Further, for each $\tau \in S_i$, let τ' and τ'' be the projection of τ to $\mathrm{vars}(\tau) \setminus E_i$ and $\mathrm{vars}(\tau) \setminus (U_i \cup E_i)$ respectively. We show that (a) $\Phi \vdash_{\mathsf{P+\forall red}} L_{\tau'}$ follows from the implicational completeness of P, and that (b) $L_{\tau''}$ can be derived from $L_{\tau'}$ by ∀-reduction in a P+∀red derivation from Φ. This completes the inductive step and the proof, since each assignment in S_{i-1} is equal to τ'' for some $\tau \in S_i$.

To show (a), let $s(\tau) = \{\sigma_1, \dots, \sigma_k\}$ be the set of assignments in S_i agreeing with τ on all blocks left of E_i. We observe that $L_{\sigma_1}, \dots, L_{\sigma_k} \vDash L_{\tau'}$, since (by definition of ∀-strategy) $s(\tau)$ contains the extension of τ' by each assignment to E_i. Hence, since each L_{σ_i} can be derived

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

[13] We assume that restriction of lines in P can be computed in polynomial time.

from Φ (by the inductive hypothesis), it follows that $\Phi \vdash_{\mathsf{P}+\forall\mathsf{red}} L_{\tau'}$ by the implicational completeness of P.

To show (b), we let β be the projection of τ to U_i. Since $\text{vars}(L_{\tau'})$ contains no variable right of U_i, in a $\mathsf{P}+\forall\mathsf{red}$ derivation from Φ we may derive $L_{\tau'}[\beta]$ from $L_{\tau'}$ by $\forall$-reduction. We observe that $B_{L_{\tau'}}|_\beta$ is the unique Boolean function on $\text{vars}(\tau'')$ that is zero only at τ''. Moreover, by restrictive closure of P, the associated Boolean function for $L_{\tau'}[\beta]$ is identical to $B_{L_{\tau'}}|_\beta$, and hence we have $L_{\tau'}[\beta] = L_{\tau''}$. □

5.3 Round-based strategy extraction in $\mathsf{P}+\forall\mathsf{red}$

The argument for soundness of $\mathsf{P}+\forall\mathsf{red}$ by strategy extraction is less straightforward than the argument for completeness. However, if we are careful enough with the details, we obtain not only a proof of soundness, but the framework for an absolute lower bound on refutation size (we take this up in the following subsection).

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 4, in $\mathsf{QU\text{-}Res}$ the response can be determined from the penultimate clause. In the general setting of $\mathsf{P}+\forall\mathsf{red}$, 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. Strictly speaking, given a line $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$ and a total assignment α to the existential variables of L, the response map returns a total assignment to the universal variables that is guaranteed to falsify $L[\alpha]$, as long as such an assignment exists.

Definition 13 (response map). *A* response map $\mathcal{R}$ *for* $\mathsf{P}+\forall\mathsf{red}$ *is any function with domain* $\{(L,\alpha) : L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}, \alpha \in \langle \text{vars}_\exists(L)\rangle\}$ *that maps each* (L,α) *to some* $\beta \in \langle \text{vars}_\forall(L)\rangle$ *such that the following holds:*

If $B_L|_\alpha$ is zero anywhere, then it is zero at β.

The case where P is resolution gives rise to a simple response map, since falsifying the universal literals is guaranteed to falsify a clause under an existential restriction, whenever possible. Hence, we simply take $\mathcal{R}(C,\alpha) = \{\neg l : l \in C_\forall\}$, which demonstrates that the response need not even depend on the existential restriction α. This is not the case for $\mathsf{P}+\forall\mathsf{red}$ in general, where the expressive capacity of lines may be much greater than that of the disjunctive Boolean functions associated with clauses.

Given an arbitrary response map for $\mathsf{P}+\forall\mathsf{red}$, we can define an algorithm whose input is a refutation of a QBF Φ and whose output is a winning $\forall$-strategy for that QBF. In a nutshell, the algorithm works 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 non-tautological line containing no existential variables.

We proceed to define *round-based strategy extraction* for $\mathsf{P}+\forall\mathsf{red}$. For the purpose of counting universal reduction steps in the following subsection, we introduce the term *reduction line* to refer to any line in π that is the antecedent in an application of universal reduction. For convenience, we also consider the conclusion of π to be a reduction line. For each response map $\mathcal{R}$ for P and each $\mathsf{P}+\forall\mathsf{red}$ refutation π of a QBF $\Phi = \exists E_1 \forall U_1 \cdots \exists E_n \forall U_n \cdot \phi$, we define a set of functions

$$\beta_i^{\mathcal{R},\pi} : \langle \text{vars}_\exists(\Phi)\rangle \to \langle U_i\rangle, \quad \text{for } i \in [n],$$

that capture the $\forall$-player's responses extracted from π. We define these functions inductively as follows: First, let $\alpha \in \langle \text{vars}_\exists(\Phi)\rangle$, and let $\beta_0(\alpha) = \emptyset$. Then, for each $i \in [n]$,

(a) let α_i be the projection of α to E_i.

(b) let Z_i be the first reduction line in π for which
  (i) $\mathrm{vars}_\exists(Z_i) \subseteq \bigcup_{j=1}^{i} E_j$, and
  (ii) $Z_i[\bigcup_{j=1}^{i} \alpha_j \cup \beta_{j-1}^{\mathcal{R},\pi}(\alpha)]$ is not a tautology.
(c) let $\beta_i^{\mathcal{R},\pi}(\alpha)$ be the projection of $\mathcal{R}(Z_i, \sigma_i)$ to U_i, where σ_i is the projection of α to $\mathrm{vars}_\exists(Z_i)$.

Note that the reduction line Z_i always exists, since the conclusion of π is $\bot$, $\mathrm{vars}(\bot) = \emptyset$, and $\bot[\tau] = \bot$ for any assignment τ (by the restrictive closure of P). The extracted strategy is constructed by taking the union of the round-by-round responses for a given existential assignment.

Definition 14 (round-based strategy extraction). *Let $\mathcal{R}$ be a response map for a base system P, and let π be a P+$\forall$red refutation of a QBF $\Phi = \exists E_1 \forall U_1 \cdots \exists E_n \forall U_n \cdot \phi$. The* (round-based) extracted strategy *for π with respect to $\mathcal{R}$ is*

$$\mathcal{S}_{\mathcal{R}}(\Phi, \pi) : \langle \mathrm{vars}_\exists(\Phi) \rangle \to \langle \mathrm{vars}_\forall(\Phi) \rangle$$
$$\alpha \mapsto \bigcup_{i=1}^{n} \beta_i^{\mathcal{R},\pi}(\alpha) \, .$$

Since it is clear that the strategy extraction algorithm terminates (each Z_i exists), we need only prove its correctness; that is, a winning $\forall$-strategy is indeed returned. In the following proof, where the choice of $\mathcal{R}$ and π is clear, we omit the superscripts and use the function symbols β_i.

It is worth noting that we do *not* employ a generalisation of Proposition 3; that is, we do not prove that a P+$\forall$red refutation is preserved by a per-line existential restriction, as is the case in QU-Res. Indeed, in the general setting, this proposition does not hold (extra propositional inferences may need to be inserted to obtain a concrete P+$\forall$red refutation). Instead, the correctness of the strategy extraction algorithm is proved by analysis of the associated Boolean functions of successive restrictions.

Lemma 15. *Let $\mathcal{R}$ be a response map for a base system P, and let π be a P+$\forall$red refutation of a QBF Φ. Then the extracted strategy for π with respect to $\mathcal{R}$ is a winning $\forall$-strategy for Φ.*

Proof. Let $\Phi = \mathcal{Q} \cdot \phi$, where $\mathcal{Q} = \exists E_1 \forall U_1 \cdots \exists E_n \forall U_n$, and let S be the extracted strategy for $\pi = L_1, \dots, L_m$ with respect to $\mathcal{R}$.

By construction, S is a $\forall$-strategy for Φ; that is, for each $\alpha, \alpha' \in \langle \mathrm{vars}_\exists(\Phi) \rangle$, if α and α' agree on the first i existential blocks, then $S(\alpha)$ and $S(\alpha')$ agree on the first i universal blocks. To see this, observe that the projection of both $S(\alpha)$ and $S(\alpha')$ to $\bigcup_{j=1}^{i} U_j$ is $\bigcup_{j=1}^{i} \beta_j(\alpha)$.

It remains to show that S is winning; that is, for each $\alpha \in \langle \mathrm{vars}_\exists(\Phi) \rangle$, $\alpha \cup S(\alpha)$ falsifies ϕ. To that end, let $\alpha \in \langle \mathrm{vars}_\exists(\Phi) \rangle$, and let $\tau_i = \bigcup_{j=1}^{i} (\alpha_j \cup \beta_j(\alpha))$ for each $i \in [n]$. Further, for each line L in π, let $\mathcal{L}_L$ be the set of lines in the subderivation of L. By induction on $i \in \{0, \dots, n\}$, we prove the following two invariants for each L in the subderivation of Z_i:

(1) If L was derived using an inference rule of P, then $\{L'[\tau_i] : L' \in \mathcal{L}_L\} \vDash L[\tau_i]$;
(2) If L was derived by $\forall$-reduction, then $B_{L[\tau_i]} = B_{L'[\tau_i]}|_\beta$ for some $L' \in \mathcal{L}_L$, where β is a partial assignment to $\mathrm{vars}_\forall(\Phi)$ with $\mathrm{vars}_\exists(L'[\tau_i]) <_{\mathcal{Q}} \mathrm{vars}(\beta)$.

At the final step $i = n$ we prove that S is indeed winning. To see this, note that τ_n is a total assignment to $\mathrm{vars}(\Phi)$, and hence $L_j[\tau_n] \in \{\top, \bot\}$ for each $j \in [m]$. Aiming for a contradiction, suppose that $A[\tau_n] = \top$ for each axiom A of π. By invariants (1) and (2), it follows that $L[\tau_n] = \top$ for each L in the subderivation of Z_n. However, we reach a contradiction, since $Z_n[\tau_n] = \bot$. To see this, recall that $Z_n[\bigcup_{j=1}^{n} \alpha_j \cup \beta_{j-1}^{\mathcal{R},\pi}]$ is not a tautology, and must therefore be falsified by $\beta_n(\alpha)$, by the definition of response map. It follows that there is some axiom A in π with $A[\tau_n] = \bot$, and hence, by the axiomatic equivalence of P, $\phi[\tau_n] = \phi[\alpha \cup S(\alpha)] = \bot$.

We prove the invariants in turn. For the base cases $i = 0$, we define Z_0 to be the conclusion of π, and τ_0 to be the empty assignment, so that the subderivation of Z_0 is π itself, and $L[\tau_0] = L$ for each L in π.

For invariant (1), let L be a line in the subderivation of Z_i that was derived using an inference rule of P. The base case is established trivially, since $\mathcal{L}_L \vDash L$ by the logical correctness of P. For the inductive step, let $i \in [n]$. Then $\{L'[\tau_{i-1}] : L' \in \mathcal{L}_L\} \vDash L[\tau_{i-1}]$, by the inductive hypothesis. Equivalently, we may say that $B_{L[\tau_{i-1}]}$ is equal to 1 wherever every function $B_{L'[\tau_{i-1}]}$ is equal to 1. Moreover, for each $L' \in \mathcal{L}_L$, we observe that

$$B_{L'[\tau_i]} = B_{L'[\tau_{i-1}][\alpha_i \cup \beta_i(\alpha)]} = B_{L'[\tau_{i-1}]}|_{\alpha_i \cup \beta_i(\alpha)},$$

by Proposition 8 and the restrictive closure of P. It follows that $B_{L[\tau_i]}$ is equal to 1 wherever every function $B_{L'[\tau_i]}$ is equal to 1. Therefore $\{L'[\tau_i] : L' \in \mathcal{L}_L\} \vDash L[\tau_i]$, and invariant (1) holds.

For invariant (2), let $L = L'[\beta]$ be a line in the subderivation of Z_i that was derived from L' by $\forall$-reduction. The base case $i = 0$ follows from the definition of $\forall$-reduction and the restrictive closure of P. For the inductive step, let $i \in [n]$, and consider two cases:

(i) Suppose that $\mathrm{vars}(L')$ contains a variable right of E_i. Then $\mathrm{vars}(\beta)$ contains only variables right of U_i, so τ_i and β are variable-disjoint assignments. It follows that

$$B_{L[\tau_i]} = B_{L'[\beta][\tau_i]} = B_{L'[\tau_i][\beta]} = B_{L'[\tau_i]}|_{\beta},$$

by Proposition 8 and restrictive closure.

(ii) On the other hand, suppose that $\mathrm{vars}_\exists(L') \subseteq \bigcup_{j=1}^{i} E_j$. Then $L'[\tau_{i-1} \cup \alpha_i]$ must be a tautology, for otherwise L' is Z_i,[14] contradicting the assumption that L is in the subderivation of Z_i. Letting $\beta' = \{l \in \beta : U_i <_{\mathcal{Q}} \mathrm{var}(l)\}$, we show that $B_{L[\tau_i]} = B_{L'[\tau_i]}|_{\beta'}$, proving the invariant. Observe that the associated Boolean function for $L'[\tau_{i-1} \cup \alpha_i]$ is everywhere 1, and hence

$$B_{L[\tau_{i-1} \cup \alpha_i]} = B_{L[\tau_{i-1}]}|_{\alpha_i} = (B_{L'[\tau_{i-1}]}|_{\beta})|_{\alpha_i} = B_{L'[\tau_{i-1} \cup \alpha_i]}|_{\beta}$$

is also everywhere 1, by Proposition 8 and restrictive closure. Similarly, both functions

$$B_{L[\tau_i]} = B_{L[\tau_{i-1} \cup \alpha_i]}|_{\beta_i(\alpha)} \quad \text{and} \quad B_{L'[\tau_i]}|_{\beta'} = B_{L'[\tau_{i-1} \cup \alpha_i]}|_{\beta_i(\alpha) \cup \beta'}$$

are everywhere 1; in fact, they are identical, since $\mathrm{vars}(L[\tau_i]) = \mathrm{vars}(L'[\tau_i]) \setminus \mathrm{vars}(\beta')$. □

If a strategy can be extracted from a $\mathsf{P}+\forall\mathsf{red}$ refutation, then the calculus is sound, hence the following corollary.

Corollary 16. *If* P *is a base system, then* $\mathsf{P}+\forall\mathsf{red}$ *is sound.*

5.4 $\mathsf{P}+\forall\mathsf{red}$ lower bounds with the Size-Cost-Capacity Theorem

In this subsection, we use the machinery of round-based strategy extraction to prove the central theorem for our lower bound technique. In Section 4, we proved the hardness of the equality formulas $\Theta(n)$ in $\mathsf{QU\text{-}Res}$ by appealing to the minimum cardinality of the range of a winning strategy. For $\Theta(n)$, the minimum cardinality is trivial to compute because the winning strategy per instance is unique. Moreover, the direct proof of the lower bound was aided by the fact the equality formulas have a Σ_3 prefix, so that all universal variables appear in a single block.

[14] Note that $\tau_i \cup \alpha_i = \bigcup_{j=1}^{i}(\alpha_j \cup \beta_{j-1}(\alpha))$, since $\beta_0(\alpha) = \emptyset$.

In order to generalise that proof method to arbitrary instances, we require a more sophisticated measure that accounts for the multiple responses collected during round-based strategy extraction. 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 17 (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 $\alpha \in \langle \mathrm{vars}_\exists(\Phi) \rangle$ to the projection of $S(\alpha)$ to U_i, and let $\mu(S) = \max\{|\mathrm{rng}(S_i)| : i \in [n]\}$. The* cost *of Φ is*

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

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

Proposition 18. *The cost of the n^{th} equality formula is 2^n.*

In the present work, the important feature of round-based strategy extraction, for the purpose of obtaining lower bounds based on cost, is that each extracted response $\beta_i(\alpha)$ can be associated with a line in the original refutation. The direct proof of hardness for $\Theta(n)$ in QU-Res used the fact that each line can contribute *at most one response*. This, of course, does not hold for P+$\forall$red in general; however, any upper bound on the number of responses per line will yield some lower bound on refutation size.

To that end, we define the concept of a *response set* for a line $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$, which is simply the set of elements $\mathcal{R}(L, \alpha)$ over all $\alpha \in \langle \mathrm{vars}_\exists(\Phi) \rangle$ for some response map $\mathcal{R}$ for P+$\forall$red.

Definition 19 (response set). *Let $\mathcal{R}$ be a response map for a base system* P*, and let $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$. The set $\{\mathcal{R}(L, \alpha) : \alpha \in \langle \mathrm{vars}_\exists(L) \rangle\}$ is a* response set *for L.*

As we defined round-based strategy extraction relative to an arbitrary response map, we may select one that minimises the size of the response sets for the lines of $\mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$, and the algorithm will still return a winning $\forall$-strategy, by Lemma 15. By selecting such a minimal response map $\mathcal{R}$, we will therefore limit the capacity for lines in the refutation to contribute multiple responses to the extracted strategy. Therefore, we associate with each P derivation the maximum number of responses that can be extracted from a single line in that derivation, with respect to a minimal response map. This notion, which we call *capacity*, captures the best-case upper bound we can place on the number of responses contributed per line.

Definition 20 (capacity). *Let* P *be a base system, let $\pi = L_1, \dots, L_m$ be a* P+$\forall$red *derivation, and let $\mu(L_i) = \min\{|R| : R \textit{ is a response set for } L_i\}$, for each $i \in [m]$. The* capacity *of π is given by*

$$\mathrm{capacity}(\pi) = \max\{\mu(L_i) : i \in [m]\}\,.$$

We saw earlier that resolution admits a response map for which the response to a given line does not depend on the existential assignment α. As such, the minimum size of a response set for any clause is 1; hence, if π is a resolution derivation, then $\mathrm{capacity}(\pi) = 1$.

Proposition 21. *Every* QU-Res *derivation has capacity equal to* 1*.*

This fact, which demonstrates the lack of expressive power of clauses, states that there exists an injection from extracted responses to lines in QU-Res refutations. Close inspection reveals that this is exactly the notion we used in the direct proof of hardness for the equality formulas (Theorem 5) in the previous section.

With definitions of cost and capacity in hand, it remains to state and prove the relationship that yields a genuine refutation-size lower bound.

Lemma 22. *Let* P *be a base system, and let* π *be a* $\mathsf{P}+\forall\mathsf{red}$ *refutation of a QBF* Φ*. Then* π *contains at least* $\mathrm{cost}(\Phi)/\mathrm{capacity}(\pi)$ *reduction lines.*

Proof. Let $\mathcal{R}$ be a response map for P with the following property: For each $L \in \mathcal{L}_{\mathsf{P}+\forall\mathsf{red}}$,

$$|\{\mathcal{R}(L,\alpha) : \alpha \in \langle \mathrm{vars}_\exists(L)\rangle\}| = \min\{|R| : R \text{ is a response set for } L\}.$$

Further, let $\Phi = \forall U_1 \exists E_1 \cdots \forall U_n \exists E_n \cdot \phi$ and let S be the extracted strategy for π with respect to $\mathcal{R}$.

Now, as in Definition 17, for each $i \in [n]$, let S_i be the function that maps each $\alpha \in \langle \mathrm{vars}_\exists(\Phi)\rangle$ to the projection of $S(\alpha)$ to U_i. Also, let $j \in [n]$ such that $|\mathrm{rng}(S_j)| = \max\{|\mathrm{rng}(S_i)| : i \in [n]\}$, and observe that $|\mathrm{rng}(S_j)| \geq \mathrm{cost}(\Phi)$.

By construction of S, $\mathrm{rng}(S_j) = \mathrm{rng}(\beta_j^{\mathcal{R},\pi})$, and each element of $\mathrm{rng}(\beta_j^{\mathcal{R},\pi})$ is the projection to U_i of $\mathcal{R}(Z_i,\alpha)$ for some reduction line Z_i in π and some $\alpha \in \langle \mathrm{vars}_\exists(\Phi)\rangle$. For a fixed line L in π, there can be at most $\mathrm{capacity}(\pi)$ elements in the set

$$\{\beta \in \mathrm{rng}(\beta_j^{\mathcal{R},\pi}) : \beta \text{ is the projection of } \mathcal{R}(L,\alpha) \text{ to } U_i \text{ for some } \alpha \in \langle \mathrm{vars}_\exists(\Phi)\rangle\}.$$

It follows that there are at least $|\mathrm{rng}(\beta_j^{\mathcal{R},\pi})|/\mathrm{capacity}(\pi)$ reduction lines in π. We hence prove the theorem, since $|\mathrm{rng}(\beta_j^{\mathcal{R},\pi})| = |\mathrm{rng}(S_j)| \geq \mathrm{cost}(\Phi)$. ☐

Lemma 22 and its proof show formally that our technique works by counting the number of universal reduction steps in the proof. Our main theorem is an immediate consequence, and hence all results proved by application of Size-Cost-Capacity are genuine QBF lower bounds.

Theorem 23 (Size-Cost-Capacity Theorem). *Let* P *be a base system, and let* π *be a* $\mathsf{P}+\forall\mathsf{red}$ *refutation of a QBF* Φ*. Then*

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

Since QU-Res derivations have unit capacity (Proposition 21), the Size-Cost-Capacity Theorem tells us that *cost alone* is an absolute refutation-size lower bound in that system.

Corollary 24. *Let* π *be a* QU-Res *refutation of a QBF* Φ*. Then* $|\pi| \geq \mathrm{cost}(\Phi)$*.*

As a first application of Size-Cost-Capacity, we therefore obtain a simple proof of the hardness of the equality formulas in QU-Res (Theorem 5), as a direct consequence of their exponential cost (Proposition 18).

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 simulate resolution, and similarly the QBF proof systems obtained with the addition of the $\forall$-reduction rule both simulate QU-Res. Our capacity upper bound for Polynomial Calculus 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 proof with large capacity.

6.1 Cutting Planes

The first proof system we analyse is Cutting Planes [24] and its extension to QBFs, $\mathsf{CP}+\forall\mathsf{red}$ [11]. Inspired by integer linear programming, Cutting Planes translates a CNF into an equivalent system of linear inequalities, and from these derives the contradiction $0 \geq 1$. 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. However, for our purposes we focus only on its use as a proof system for unsatisfiable CNFs.

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 25 (Cutting Planes [24]). *A line L in a Cutting Planes* (CP) *proof is a linear inequality $a_1x_1 + \ldots a_nx_n \geq A$ where* $\mathrm{vars}(L) = \{x_1, \ldots, x_n\}$ *and $a_1, \ldots, a_n, A \in \mathbb{Z}$.*[15]

A Cutting Planes derivation of a line $L \in \mathcal{L}_{\mathsf{CP}}$ from a CNF ϕ consists of a sequence of lines $L_1, \ldots, L_m$ where $L_m = L$ and each line $L_i \in \mathcal{L}_{\mathsf{CP}}$ 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:	$\dfrac{}{\sum_{l \in C} R(l) \geq 1}$	for any clause $C \in \phi$, where $R(x) = x$, $R(\neg x) = 1 - x$
Boolean Axiom:	$\dfrac{}{x \geq 0}$ $\qquad$ $\dfrac{}{-x \geq -1}$	for any variable x
Linear combination:	$\dfrac{\sum_i a_i x_i \geq A \qquad \sum_i b_i x_i \geq B}{\sum_i (\alpha a_i + \beta b_i) x_i \geq \alpha A + \beta B}$	for any $\alpha, \beta \in \mathbb{N}$
Division:	$\dfrac{\sum_i c a_i x_i \geq A}{\sum_i a_i x_i \geq \lceil \frac{A}{c} \rceil}$	for any non-zero $c \in \mathbb{N}$

Fig. 4. The rules of the Cutting Planes proof system (CP)

It is straightforward to see that Cutting Planes p-simulates resolution. 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 [37]. The same argument shows that $\mathsf{CP}+\forall\mathsf{red}$ is exponentially stronger than $\mathsf{QU\text{-}Res}$.

Despite the apparent strength of $\mathsf{CP}+\forall\mathsf{red}$ compared to $\mathsf{QU\text{-}Res}$, any proof in $\mathsf{CP}+\forall\mathsf{red}$ still only has unit capacity. This comes about as the left hand side of any inequality L is simply a linear combination of variables. Any response evaluates the term b_iu_i either to 0 or to b_i. By assigning the values of $u_1, \ldots, u_n$ according to the sign of their coefficients, there is a response such that each term evaluates to $\min\{0, b_i\}$. This universal response minimises the left hand side for any existential assignment, and so forms a response set of cardinality 1.

Proposition 26. *For every* $\mathsf{CP}+\forall\mathsf{red}$ *derivation π,* $\mathrm{capacity}(\pi) = 1$.

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

Proof. We show that for any line $L \in \mathcal{L}_{\mathsf{CP+\forall red}}$, there is a response set of cardinality 1. It follows that capacity$(\pi) = 1$ as for every $L \in \pi$, $L \in \mathcal{L}_{\mathsf{CP+\forall red}}$ since π is a CP+∀red proof, so every line $L \in \pi$ has a response set of cardinality 1.

Let $L \in \mathcal{L}_{\mathsf{CP+\forall red}}$, with $\mathrm{vars}_\exists(L) = \{x_1, \dots, x_m\}$ and $\mathrm{vars}_\forall(L) = \{u_1, \dots, u_n\}$. Then L is of the form

$$a_1x_1 + \cdots + a_mx_m + b_1u_1 + \cdots + b_nu_n \geq C$$

for some constants $a_i, b_j, C \in \mathbb{Z}$.

Define the response $\gamma_L \in \langle \mathrm{vars}_\forall(L) \rangle$ which assigns to each variable u_i the value $\mathrm{sgn}(b_i)$, where $\mathrm{sgn} : \mathbb{Z} \to \{0, 1\}$ with $\mathrm{sgn}(x) = 0$ if and only if $x \geq 0$. This assignment includes in the sum only those b_i which are negative, thus minimising the value of the sum for any existential assignment $\alpha \in \langle \mathrm{var}_\exists(L) \rangle$. It remains to show that the map $\mathcal{R}$ where $\mathcal{R}(L, \alpha) = \gamma_L$ for all $L \in \mathcal{L}_{\mathsf{CP+\forall red}}$ and $\alpha \in \langle \mathrm{vars}_\exists(L) \rangle$ is a response map, and so for each line $L \in \pi$, $\{\gamma_L\}$ is a response set for L.

Fix an assignment α to $\mathrm{vars}_\exists(L)$; the restricted line $L[\alpha]$ is $\sum_{i=1}^m a_i\alpha(x_i) + b_1u_1 + \cdots + b_nu_n \geq C$. Letting the constant $D = C - \sum_{i=1}^m a_i\alpha(x_i)$, we see that $B_L|_\alpha = B_{L[\alpha]}$ is zero if and only if $\sum_{j=1}^n b_ju_j < D$. By the definition of γ_L, $\sum_{j=1}^n b_j\gamma_L(u_j) = \sum_{j:b_j<0} b_j$, so $\sum_{j=1}^n b_j\gamma_L(u_j) \leq \sum_{j=1}^n b_j\beta(u_j)$ for all $\beta \in \langle \mathrm{vars}_\forall(L) \rangle$.

Suppose that for some $\beta \in \langle \mathrm{vars}_\forall(L) \rangle$ we have $B_L|_\alpha(\beta) = 0$. By the equivalence above, it must be the case that $\sum_{j=1}^n b_j\beta(u_j) < D$. The assignment γ_L is such that $\sum_{j=1}^n b_j\gamma_L(u_j) \leq \sum_{j=1}^n b_j\beta(u_j) < D$, and so $B_L|_\alpha(\gamma_L) = 0$. We conclude that $\mathcal{R}$ is a response map for CP+∀red, and that capacity$(\pi) = 1$ for any CP+∀red proof π ☐

Since we have determined the capacity of any CP+∀red proof π, we can apply the Size-Cost-Capacity Theorem (Theorem 23). This yields a lower bound on the size of CP+∀red proofs of a QBF Φ determined solely by the cost of Φ.

Corollary 27. *Let π be a* CP+∀red *refutation of a QBF Φ. Then $|\pi| \geq \mathrm{cost}(\Phi)$.*

Hence, even in the stronger system of CP+∀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 [21] presents an algebraic approach to proving unsatisfiability. A CNF ϕ is translated into a set of polynomials, for which any assignment where all polynomials evaluate to zero corresponds to a satisfying assignment for ϕ, and vice versa. Replacing the axiom rules with any finite set of polynomials over $\mathbb{Q}$, Polynomial Calculus is complete for any set of polynomials without a common zero, but here we consider Polynomial Calculus only as a proof system for unsatisfiable CNFs.

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, Polynomial Calculus works with 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 $V(\neg x) = (1 - x)$.[16] We also include the axioms $x^2 - x = 0$ to ensure only Boolean solutions.

[16] Assigning 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 feature of algebraic proof systems.

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 $\neg x_1 \vee \cdots \vee \neg x_n$ 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 [1], 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 28 (Polynomial Calculus with Resolution [1]). *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 $L \in \mathcal{L}_{\mathsf{PCR}}$ from a CNF ϕ is a sequence of lines $L_1, \dots, L_m$ in $\mathcal{L}_{\mathsf{PCR}}$ 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 $1 = 0$.*

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

Fig. 5. The rules of Polynomial Calculus with Resolution (PCR)

As Polynomial Calculus with Resolution is clearly strictly stronger than Polynomial Calculus, 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 CP+∀red, not all proofs in PCR+∀red have unit capacity. This can be seen by considering 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 $u = 1 - x$. Both possible responses are necessary in any response set for this line, so if a PCR+∀red proof π contains such a line, then $\text{capacity}(\pi) \geq 2$.

While the size of response sets required for lines in $\mathcal{L}_{\mathsf{PCR+\forall red}}$ is in fact unbounded, the size of a proof in PCR+∀red is not measured by the number of lines in the proof, but by the number of monomials in the proof. It is thus sufficient to upper bound the size of a response set for a line by the number of monomials in that line.

Proposition 29. *If π is a* PCR+∀red *proof where each line contains at most M monomials, then* $\text{capacity}(\pi) \leq M$*.*

To prove this bound, the key observation is that the important feature for determining whether a response is winning on a line $L \in \mathcal{L}_{\mathsf{PCR+\forall red}}$ is the evaluation of the response on the distinct monomials of L, rather than the assignment to the individual variables. Rather than considering responses, we therefore consider the vectors of the values the responses take on the distinct monomials in L. If two responses correspond to the same vector, these responses are winning for precisely the same existential assignments, so only one need be in a response set. For any given vector, at most one of the responses mapping to it will be in a minimal response set for L, so an upper bound on the size of a winning set of these vectors is sufficient to upper bound the size of a response set for L.

If, for some existential assignment, a set of these vectors does not contain a winning 'response', we show that any vector in their span is not a winning response either. We can therefore find a response set for which the corresponding vectors are linearly independent. As these vectors have dimension at most M, this provides the upper bound.

Proof (of Proposition 29). To upper bound capacity(π), it is enough to find, for any line $L \in \pi$, a response set of cardinality at most M. To this end, fix a line $L \in \pi$ 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 $\mathrm{vars}_\exists(L)$ and the v_i are distinct monomials in $\mathrm{vars}_\forall(L)$. Denote by $f_j[\alpha]$ and $v_j[\beta]$ the values obtained by evaluating f_j, respectively v_j, according to the assignments $\alpha \in \langle \mathrm{vars}_\exists(L)\rangle$, respectively $\beta \in \langle \mathrm{vars}_\forall(L)\rangle$.

Observe that since L contains at most M monomials, $N \leq M$. We construct a (partial) response map $\mathcal{R}_L : \langle \mathrm{vars}_\exists(L)\rangle \to \langle \mathrm{vars}_\forall(L)\rangle$ for which the corresponding response set for L has cardinality at most N. This generates a response set of size at most M for each line $L \in \pi$, and so capacity(π) $\leq M$.

Enumerate the elements of $\langle \mathrm{vars}_\exists(L)\rangle$ as $\langle \mathrm{vars}_\exists(L)\rangle = \{\alpha_1, \dots, \alpha_m\}$. We construct a sequence of functions $\mathcal{R}_L^i : \{\alpha_1, \dots, \alpha_i\} \to \langle \mathrm{vars}_\forall(L)\rangle$ such that $\mathcal{R}_L^0$ is the empty function, $\mathcal{R}_L^i$ extends $\mathcal{R}_L^{i-1}$, and $\mathcal{R}_L^m$ is the partial response map $\mathcal{R}_L$ above. Moreover, for each $0 \leq i \leq m$, $|\mathrm{rng}(\mathcal{R}_L^i)| \leq N$, in particular $|\mathrm{rng}(\mathcal{R}_L^m)| \leq N$.

The construction of the function $\mathcal{R}_L^m$ involves, in effect, going through each of the assignments $\alpha \in \langle \mathrm{vars}_\exists(L)\rangle$ in turn and choosing a response. At each stage, if there is a suitable response that has been chosen before, we choose it again as we are aiming for a response set of minimal size. If there is no suitable previously chosen response, we show that from the newly chosen response and the set of previously chosen responses, there is an injection into a linearly independent subset of $\mathbb{Q}^N$.

First, if $B_L|_\alpha$ is everywhere 1 for all $\alpha \in \langle \mathrm{vars}_\exists(L)\rangle$, then $L[\alpha]$ is a tautology. In this case, all functions $\langle \mathrm{vars}_\exists(L)\rangle \to \langle \mathrm{vars}_\forall(L)\rangle$ are suitable partial response maps. Define $\mathcal{R}_L^i(\alpha) = \mathbf{0}$ for all $1 \leq i \leq M$ and all $\alpha \in \langle \mathrm{vars}_\exists(L)\rangle$, whence $|\mathrm{rng}(\mathcal{R}_L^m)| = 1 \leq M$. Otherwise, we assume without loss of generality that $B_L|_{\alpha_1}$ is not everywhere 1.

Let $R_i := \mathrm{rng}(\mathcal{R}_L^i)$. For any $\beta \in \langle \mathrm{vars}_\forall(L)\rangle$, denote by $\boldsymbol{v}[\beta]$ the vector $(v_1[\beta], \dots, v_N[\beta]) \in \{0,1\}^N$. Since the value of $B_L|_\alpha(\beta)$ is determined by the values of $v_j[\beta]$, it is clear in the construction below that $\boldsymbol{v}[\beta] \neq \boldsymbol{v}[\beta']$ for any distinct $\beta, \beta' \in R_i$. In constructing the functions $\mathcal{R}_L^i$, we show inductively that the set $V_i = \{\boldsymbol{v}[\beta] : \beta \in R_i\}$ is linearly independent (as a subset of $\mathbb{Q}^N$). Since $V_i \subseteq \{0,1\}^N$, and $|V_i| = |R_i|$, this provides the upper bound on $|R_i|$.

The conditions above are clearly true for the empty function $\mathcal{R}_L^0$. Given a function $\mathcal{R}_L^{i-1}$, we define $\mathcal{R}_L^i$ as follows:

- If $B_L|_{\alpha_i}$ is everywhere 1, i.e. $\sum_{j=1}^N f_j[\alpha_i] v_j[\beta] = 0$ for all $\beta \in \langle \mathrm{vars}_\forall(L)\rangle$, any response can be chosen for $\mathcal{R}_L^i(\alpha_i)$. 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 must be linearly independent.
- If $B_L|_{\alpha_i}(\beta) = 0$ 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 $B_L|_{\alpha_i}$ is not everywhere 1, but $B_L|_{\alpha_i}(\beta) = 1$ 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 $B_L|_{\alpha_i}(\beta_{k+1}) = 0$, since $B_L|_{\alpha_i}$ is not

everywhere 1. Define $\mathcal{R}_L^i(\alpha_i) = \beta_{k+1}$, with $R_i = \{\beta_1, \ldots, \beta_k, \beta_{k+1}\}$. We now need only show that $V_i = \{\boldsymbol{v}[\beta_j] : 1 \leq j \leq k+1\}$ is linearly independent.
For any $1 \leq l \leq k$, $B_L|_{\alpha_i}(\beta_{k+1}) \neq B_L|_{\alpha_i}(\beta_l)$, so $\boldsymbol{v}[\beta_{k+1}] \neq \boldsymbol{v}[\beta_l]$. Suppose there is some linear dependence relation on V_i. Since V_{i-1} is linearly independent, $\boldsymbol{v}[\beta_{k+1}]$ must have a non-zero coefficient in any such linear combination, hence there are constants $c_1, \ldots, c_k \in \mathbb{Q}$ such that $\sum_{t=1}^{k} c_t \boldsymbol{v}[\beta_t] = \boldsymbol{v}[\beta_{k+1}]$. 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 $\mathbb{Q}^N$, so has cardinality at most N. Since $|R_m| = |V_m|$, the map $\mathcal{R}_L = \mathcal{R}_L^m$ satisfies $|\text{rng}(\mathcal{R}_L)| \leq N \leq M$. Defining $\mathcal{R}(L, \alpha) = \mathcal{R}_L(\alpha)$, we obtain a response map which produces a response set of size at most M for any line $L \in \pi$. We conclude that capacity$(\pi) \leq M$. ◻

The effect of this bound is to show that PCR+∀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+∀red proofs of a QBF Φ based solely on cost(Φ), since small proofs also have small capacity.

Corollary 30. *Let π be a* PCR+∀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 29, capacity$(\pi) \leq |\pi|$. Applying Size-Cost-Capacity (Theorem 23), we conclude that $|\pi| \geq \frac{\text{cost}(\Phi)}{|\pi|}$, i.e. $|\pi| \geq \sqrt{\text{cost}(\Phi)}$. ◻

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

A Frege proof system is a 'textbook' propositional proof system, in which lines are arbitrary formulas in propositional variables, the constants $\top, \bot$ and the connectives $\wedge, \vee, \neg$. The rules of a Frege system comprise a set of axiom schemes and inference rules, which must be implicationally complete [23]; all such systems are equivalent [23, 43].

Proposition 31. *There is a* Frege+∀red *refutation of the n^{th} equality formula $\Theta(n)$ 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 $t_1 \vee \cdots \vee t_n$ 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 $K_0 = \bot$. 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 $\Theta(n)$. ◻

Let π be the Frege+$\forall$red proof described above. By the Size-Cost-Capacity Theorem (Theorem 23), we know that capacity(π) is of exponential size. We can show this directly by considering the 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 $\langle\{x_1, \dots, x_n\}\rangle$, any response set for L must have size 2^n, despite L being a Frege+$\forall$red line of size polynomial in 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 use this new lower bound technique to construct a class of randomly generated QBFs which with high probability are false and have large cost. 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.

We begin by defining a class of random formulas, for which we show a cost lower bound for appropriate values of the parameters m and c.

Definition 32. *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 [19,25]. 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.

While the QBFs $Q(n, m, c)$ are randomly generated, they still have a structure which enables us to better understand them by studying simpler randomly generated QBFs. The QBFs we look at for this purpose are $\Psi_i := \exists Y_i \forall X_i \cdot \bigwedge_{j=1}^{cn} C_i^j$, which are generated using the same clauses as the (1,2)-QCNF model (Definition 35), which is a common model of random QBFs [25].

The equivalence of $Q(n,m,c)$ to $\bigvee_{i=1}^{n} \Psi_i$ is clear. Moreover, any winning $\forall$-strategy for $Q(n,m,c)$ is a winning $\forall$-strategy for each individual Ψ_i simultaneously. In order to show that, for suitable values of m and c, $Q(n,m,c)$ is false and has large cost, we first show that, with high probability, all the Ψ_i are false. We then further show that with probability $1-o(1)$, a linear number of the Ψ_i require at least 2 different responses in $\langle X_i \rangle$. From this we obtain a lower bound on $\mathrm{cost}(Q(n,m,c))$.

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 [20], Goerdt [35] and de la Vega [57], 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 [58], which is necessary for our present work.

Theorem 33 (de la Vega [58]). *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 $1-o\left(n^{-1}\right)$.*

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 33; 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 33. The possibility of repeating an existential clause many times with different universal literals makes this non-trivial, but the proof is relatively straightforward.

Lemma 34. *For each $1 \le i \le n$, let ψ_i be a set of cn clauses picked uniformly at random from the set of clauses containing 1 literal from $X_i = \{x_i^1, \ldots, x_i^m\}$ and 2 literals from $Y_i = \{y_i^1, \ldots, y_i^n\}$. If $m \le \log_2(n)$ and $c > 1$, then with probability $1-o(1)$, $\Psi_i := \exists Y_i \forall X_i \cdot \psi_i$ is false for all $1 \le i \le n$.*

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 33, 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 variables of a clause, and $2m \le 2\log(n)$ 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, we consider choosing cn clauses at random 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

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

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 $1 - o\left(n^{-1}\right)$, by Theorem 33. The probability of Ψ_i being false is therefore at least $P(Z > kn) \cdot \left(1 - o\left(n^{-1}\right)\right) = 1 - o\left(n^{-1}\right)$.

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 $\left(1 - o\left(n^{-1}\right)\right)^n = 1 - o(1)$. □

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 $S : \langle Y_i \rangle \to \langle X_i \rangle$ 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 35 (Chen and Interian [19]). *A (1,2)-QCNF is a QBF of the form $\forall X \exists Y \cdot \phi(X,Y)$ where $X = \{x_1, \ldots, x_m\}$, $Y = \{y_1, \ldots, y_n\}$ 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 36 (Creignou et al. [25]). *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 33 and the upper bound from Theorem 36. 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 $\forall X \exists Y \cdot \psi$, whence the latter is false. This gives us the bound $\text{cost}(\Psi_i) \geq 2$ with high probability.

Lemma 37. *Let ψ be a set of cn clauses picked uniformly at random from the set of clauses containing 1 literal from the set $X = \{x_1, \ldots, x_m\}$ and 2 literals from $Y = \{y_1, \ldots, y_n\}$. If $1 < c < 2$ and $m \leq \log_2(n)$, then with high probability $\Psi := \exists Y \forall X \cdot \psi$ is false, and* $\mathrm{cost}(\Psi) \geq 2$.

Proof. Observe from the proof of Lemma 34 that, as $c > 1$, Ψ is false with probability $1 - o(n^{-1})$ and $\mathrm{cost}(\Psi) \geq 1$.

Suppose $\mathrm{cost}(\Psi) = 1$, 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 $\alpha \in \langle Y \rangle$, $\psi[\alpha][\beta] = \bot$. We can use β as a winning strategy for $\Psi' = \forall X \exists Y.\psi$, defining $S'(\emptyset) = \beta$. Since $\psi[\beta][\alpha] = \bot$ 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 36), and so such a $\beta \in \langle X \rangle$ exists with probability $o(1)$.

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

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 $\mathrm{cost}(Q(n, m, c))$.

Lemma 38. *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 $X_i = \{x_i^1, \ldots, x_i^m\}$ and 2 literals from $Y_i = \{y_i^1, \ldots, y_i^n\}$. Further suppose that $m \leq \log_2(n)$ and c, l are any constants with $1 < c < 2$, $l < 1$. With high probability $\Psi_i = \exists Y_i \forall X_i \cdot \psi_i$ is false for every $1 \leq i \leq n$ and at least ln of the Ψ_i have* $\mathrm{cost}(\Psi_i) \geq 2$.

Proof. By Lemma 37, for each Ψ_i, the probability that $\mathrm{cost}(\Psi_i) \geq 2$ 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 $\mathrm{cost}(\Psi_i) \geq 2$ is at most

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

which for sufficiently large n can be upper bounded by

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

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

Lemma 38 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 $\mathrm{cost}(\Psi_i) \geq 2$. 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 $\mathrm{cost}(Q(n, m, c))$ simply by multiplying $\mathrm{cost}(\Psi_i)$ for each $i \in [n]$, as responses on $\mathrm{vars}_\forall(\Psi_i)$ may now vary depending on the assignment of variables in some other Ψ_j. Instead, we use the fact that if $\mathrm{cost}(\Psi_i) \geq 2$, 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.

Proposition 39. *Let $1 < c < 2$ be a constant, and let $m \leq (1 - \epsilon) \log_2(n)$ for some constant $\epsilon > 0$. With probability $1 - o(1)$, $Q(n, m, c)$ is false and* $\mathrm{cost}(Q(n, m, c)) = 2^{\Omega(n^\epsilon)}$.

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

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

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 $1 \le i \le n$, 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 38, with probability $1 - o(1)$, all the Ψ_i are false, and at least ln of the Ψ_i have $\mathrm{cost}(\Psi_i) \ge 2$. It therefore suffices to show that if all the Ψ_i are false and $\mathrm{cost}(\Psi_i) \ge 2$ holds for at least ln of the Ψ_i, then $Q(n, m, c)$ is false and $\mathrm{cost}(Q(n, m, c)) \ge 2^{\Omega(n^\epsilon)}$.

If each Ψ_i is false, there is some winning strategy $S_i : \langle Y_i \rangle \to \langle X_i \rangle$ for each $i \in [n]$. Define $S : \langle Y_1, \dots, Y_n \rangle \to \langle X_1, \dots, X_n \rangle$ 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, $C_i^j[\alpha_i][S_i(\alpha_i)] = \perp$ for some $1 \le j \le cn$, 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 $\mathrm{cost}(Q(n, m, c)) \ge 2^{\Omega(n^\epsilon)}$. 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 $\Psi_1, \dots, \Psi_{ln}$, 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 $\Psi_1, \dots, \Psi_{ln}$.

Since $|X_i| \le (1 - \epsilon) \log_2(n)$, 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 $N = 2^m \le n^{(1-\epsilon)}$.

Let $B = \mathrm{rng}(S)$ for some winning $\forall$-strategy S for $Q(n, m, c)$. To lower bound $\mathrm{cost}(Q(n, m, c))$, 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. $B \subseteq \{(\beta_{j_1}^i, \dots, \beta_{j_{ln}}^{ln}) : j_1, \dots, j_{ln} \in [N]\}$. 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* $1 \le i \le ln$.

For each $1 \le i \le ln$, 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 $\mathrm{cost}(\Psi_i) = 1$. 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 $j \ne j'$, so defining A to be a multiset ensures $|A| = N^{ln}$. Given a response $\beta \in B$, we bound the size of the multiset

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

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 $1 \le i \le ln$. Since $\beta|_{X_i} = \beta_{j_i}^i$, by the definition of $\alpha^i_{j_i}$, $\alpha|_{Y_i} \neq \alpha^i_{j_i}$ for all $1 \leq i \leq ln$, as $\Psi_i[\alpha^i_{j_i}][\beta^i_{j_i}] = \top$. The restriction $\alpha|_{Y_i} \neq \alpha^i_{j_i}$ restricts the set A_β to

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

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 $|\mathrm{rng}(S)| \geq 2^{\Omega(n^\epsilon)}$ for any winning $\forall$-strategy S. There is only one block of universal variables in $Q(n,m,c)$, and so

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

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 $\mathrm{cost}(Q(n,m,c)) \geq 2^{\Omega(n^\epsilon)}$. Lemma 38 states that these conditions hold with probability $1 - o(1)$, and this completes the proof. □

Proposition 39 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 40. *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+∀red *or* PCR+∀red *refutation of $Q(n,m,c)$ requires size $2^{\Omega(n^\epsilon)}$.*

As previously, the greater capacity of lines in Frege+∀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 ∀-reduction to each clause results in an unsatisfiable 2-SAT instance, which has a linear size resolution refutation. This immediately gives a short Frege+∀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. [41] using Size-Cost-Capacity.

Definition 41 (Kleine Büning et al. [41]). *The formulas $\kappa(n)$ are defined to be*

$$\kappa(n) := \exists y_0 \exists y_1 \exists y_1' \forall u_1 \ldots \exists y_k \exists y_k' \forall u_k \ldots \exists y_n \exists y_n' \forall u_n \exists t_1 \ldots t_n \cdot \bigwedge_{i=1}^{2n} C_i \wedge C_i'$$

where the matrix contains the clauses

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

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 [10, 41], 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 [4]. 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 [4, 10, 28, 44]).

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 42. *The QBF* $\lambda'(n) := \exists y_0 y_1 y'_1 \forall u_1 \dots \exists y_n y'_n \forall u_n v_1 \dots v_n \exists t_1 \dots t_n \cdot \bigwedge_{i=1}^{2n} D_i \wedge D'_i$ *has cost* 2^n.

Proof. We consider the response of any winning $\forall$-strategy to the 2^n distinct assignments in the set $A = \{\alpha \in \langle\{y_1, y'_1, \dots, y_n, y'_n\}\rangle : \alpha(y_k) \neq \alpha(y'_k)$ 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

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

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 $|\mathrm{rng}(S_n)| = 2^n$, whence $\mathrm{cost}(\lambda'(n)) = 2^n$. ☐

We therefore obtain the following hardness result, which was known for QU-Res [4], but also lifts to CP+∀red and PCR+∀red.

Corollary 43. *Any* QU-Res*,* CP+∀red *or* PCR+∀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+∀red or PCR+∀red, which is an immediate consequence of Proposition 42 and the results of Section 6. ☐

This lower bound for QU-Res also yields the lower bound on Q-Res proofs of $\kappa(n)$, previously shown in [10, 41]. 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 [12]. 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

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

In strong proof systems such as Frege+∀red, superpolynomial proof size lower bounds can be completely characterised: they are either a propositional lower bound or a circuit lower bound [13]. All QBFs we have considered have no underlying propositional hardness, and winning ∀-strategies can be computed by small circuits, even in very restricted circuit classes. As such, all these QBFs are easy for Frege+∀red.

However, for weaker proof systems, such as QU-Res and the others we have discussed, propositional hardness and circuit lower bounds alone are not the complete picture. 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.

References

1. Alekhnovich, M., Ben-Sasson, E., Razborov, A.A., Wigderson, A.: Space complexity in propositional calculus. SIAM Journal on Computing 31(4), 1184–1211 (2002)
2. Alekhnovich, M., Razborov, A.A.: Lower bounds for polynomial calculus: Non-binomial case. In: 42nd Annual Symposium on Foundations of Computer Science, FOCS 2001. pp. 190–199 (2001)

3. Balabanov, V., Jiang, J.R.: Unified QBF certification and its applications. Formal Methods in System Design 41(1), 45–65 (2012)
4. Balabanov, V., Widl, M., Jiang, J.R.: QBF resolution systems and their proof complexities. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 154–169 (2014)
5. Beame, P., Pitassi, T.: Simplified and improved resolution lower bounds. In: Symposium on Foundations of Computer Science (FOCS). pp. 274–282 (1996)
6. Ben-Sasson, E., Impagliazzo, R.: Random CNF's are hard for the polynomial calculus. Computational Complexity 19(4), 501–519 (2010)
7. Benedetti, M., Mangassarian, H.: QBF-based formal verification: Experience and perspectives. Journal on Satisfiability, Boolean Modeling and Computation (JSAT) 5(1-4), 133–191 (2008)
8. Beyersdorff, O., Bonacina, I., Chew, L.: Lower bounds: From circuits to QBF proof systems. In: ACM Conference on Innovations in Theoretical Computer Science (ITCS). pp. 249–260 (2016)
9. Beyersdorff, O., Chew, L., Janota, M.: On unification of QBF resolution-based calculi. In: International Symposium on Mathematical Foundations of Computer Science (MFCS). pp. 81–93 (2014)
10. Beyersdorff, O., Chew, L., Janota, M.: Proof complexity of resolution-based QBF calculi. In: International Symposium on Theoretical Aspects of Computer Science (STACS). Leibniz International Proceedings in Informatics (LIPIcs), vol. 30, pp. 76–89 (2015)
11. Beyersdorff, O., Chew, L., Mahajan, M., Shukla, A.: Understanding Cutting Planes for QBFs. In: 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2016). Leibniz International Proceedings in Informatics (LIPIcs), vol. 65, pp. 40:1–40:15 (2016)
12. Beyersdorff, O., Hinde, L., Pich, J.: Reasons for hardness in QBF proof systems. Electronic Colloquium on Computational Complexity (ECCC) 24, 44 (2017)
13. Beyersdorff, O., Pich, J.: Understanding Gentzen and Frege systems for QBF. In: Symposium on Logic in Computer Science (LICS). pp. 146–155 (2016)
14. Bloem, R., Könighofer, R., Seidl, M.: SAT-based synthesis methods for safety specs. In: International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI). pp. 1–20 (2014)
15. Brummayer, R., Lonsing, F., Biere, A.: Automated testing and debugging of SAT and QBF solvers. In: Theory and Applications of Satisfiability Testing - (SAT) 2010. pp. 44–57 (2010)
16. Buss, S.R.: Towards NP-P via proof complexity and search. Ann. Pure Appl. Logic 163(7), 906–917 (2012)
17. Cashmore, M., Fox, M., Giunchiglia, E.: Partially grounded planning as quantified Boolean formula. In: International Conference on Automated Planning and Scheduling (ICAPS) (2013)
18. Chen, H.: Proof complexity modulo the polynomial hierarchy: Understanding alternation as a source of hardness. In: International Colloquium on Automata, Languages, and Programming (ICALP). pp. 94:1–94:14 (2016)
19. Chen, H., Interian, Y.: A model for generating random quantified Boolean formulas. In: International Joint Conference on Artificial Intelligence. pp. 66–71 (2005)
20. Chvátal, V., Reed, B.A.: Mick gets some (the odds are on his side). In: Symposium on Foundations of Computer Science. pp. 620–627 (1992)
21. Clegg, M., Edmonds, J., Impagliazzo, R.: Using the Groebner basis algorithm to find proofs of unsatisfiability. In: Proc. 28th ACM Symposium on Theory of Computing. pp. 174–183 (1996)
22. Cook, S.A., Nguyen, P.: Logical Foundations of Proof Complexity. Cambridge University Press (2010)
23. Cook, S.A., Reckhow, R.A.: The relative efficiency of propositional proof systems. Journal of Symbolic Logic 44(1), 36–50 (1979)
24. Cook, W., Coullard, C.R., Turán, G.: On the complexity of cutting-plane proofs. Discrete Applied Mathematics 18(1), 25–38 (1987)
25. Creignou, N., Daudé, H., Egly, U., Rossignol, R.: Exact location of the phase transition for random (1, 2)-QSAT. RAIRO - Theor. Inf. and Applic. 49(1), 23–45 (2015)
26. Dershowitz, N., Hanna, Z., Katz, J.: Space-efficient bounded model checking. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 502–518. Springer (2005)
27. Díaz, J., Kirousis, L.M., Mitsche, D., Pérez-Giménez, X.: A new upper bound for 3-SAT. In: IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2008. pp. 163–174 (2008)
28. Egly, U.: On stronger calculi for QBFs. In: Theory and Applications of Satisfiability Testing - SAT 2016 - 19th International Conference, Bordeaux, France, July 5-8, 2016, Proceedings. pp. 419–434 (2016)
29. Egly, U., Kronegger, M., Lonsing, F., Pfandler, A.: Conformant planning as a case study of incremental QBF solving. Annals of Mathematics and Artificial Intelligence 80(1), 21–45 (2017)
30. Egly, U., Lonsing, F., Widl, M.: 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). pp. 291–308 (2013)
31. Fleming, N., Pankratov, D., Pitassi, T., Robere, R.: Random CNFs are hard for cutting planes. CoRR abs/1703.02469 (2017)

32. Franco, J., Paull, M.C.: Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem. Discrete Applied Mathematics 5(1), 77–87 (1983)
33. Gelder, A.V.: Contributions to the theory of practical quantified Boolean formula solving. In: International Conference on Principles and Practice of Constraint Programming (CP). pp. 647–663 (2012)
34. Giunchiglia, E., Marin, P., Narizzano, M.: Reasoning with quantified Boolean formulas. In: Handbook of Satisfiability, pp. 761–780. IOS Press (2009)
35. Goerdt, A.: A threshold for unsatisfiability. J. Comput. Syst. Sci. 53(3), 469–486 (1996)
36. Goultiaeva, A., Gelder, A.V., Bacchus, F.: A uniform approach for generating proofs and strategies for both true and false QBF formulas. In: International Joint Conference on Artificial Intelligence IJCAI. pp. 546–553 (2011)
37. Haken, A.: The intractability of resolution. Theoretical Computer Science 39, 297–308 (1985)
38. Hrubes, P., Pudlák, P.: Random formulas, monotone circuits, and interpolation. Electronic Colloquium on Computational Complexity (ECCC) 24, 42 (2017)
39. Janota, M., Marques-Silva, J.: Expansion-based QBF solving versus Q-resolution. Theor. Comput. Sci. 577, 25–42 (2015)
40. Jeřábek, E.: Weak pigeonhole principle, and randomized computation. Ph.D. thesis, Faculty of Mathematics and Physics (2005)
41. Kleine Büning, H., Karpinski, M., Flögel, A.: Resolution for quantified Boolean formulas. Information and Computation 117(1), 12–18 (1995)
42. Kontchakov, R., Pulina, L., Sattler, U., Schneider, T., Selmer, P., Wolter, F., Zakharyaschev, M.: Minimal module extraction from DL-lite ontologies using QBF solvers. In: International Joint Conference on Artificial Intelligence (IJCAI). pp. 836–841. AAAI Press (2009)
43. Krajíček, J.: Bounded Arithmetic, Propositional Logic, and Complexity Theory, Encyclopedia of Mathematics and Its Applications, vol. 60. Cambridge University Press, Cambridge (1995)
44. Lonsing, F., Egly, U., Seidl, M.: Q-resolution with generalized axioms. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 435–452 (2016)
45. Malik, S., Zhang, L.: Boolean satisfiability from theoretical hardness to practical success. Communications of the ACM 52(8), 76–82 (2009)
46. Meel, K.S., Vardi, M.Y., Chakraborty, S., Fremont, D.J., Seshia, S.A., Fried, D., Ivrii, A., Malik, S.: Constrained sampling and counting: Universal hashing meets SAT solving. In: Beyond NP AAAI Workshop (2016)
47. Narizzano, M., Peschiera, C., Pulina, L., Tacchella, A.: Evaluating and certifying QBFs: A comparison of state-of-the-art tools. AI Communications 22(4), 191–210 (2009)
48. Peitl, T., Slivovsky, F., Szeider, S.: Long distance Q-resolution with dependency schemes. In: International Conference on Theory and Applications of Satisfiability Testing (SAT). pp. 500–518 (2016)
49. Pulina, L.: The ninth QBF solvers evaluation - preliminary report. In: Proceedings of the 4th International Workshop on Quantified Boolean Formulas (QBF 2016). pp. 1–13 (2016)
50. Razborov, A.A.: Lower bounds for the size of circuits of bounded depth with basis $\{\wedge, \oplus\}$. Math. Notes Acad. Sci. USSR 41(4), 333–338 (1987)
51. Rintanen, J.: Asymptotically optimal encodings of conformant planning in QBF. In: National Conference on Artificial Intelligence (AAAI). pp. 1045–1050. AAAI Press (2007)
52. Silva, J.P.M., Sakallah, K.A.: GRASP - a new search algorithm for satisfiability. In: International Conference on Computer-Aided Design (ICCAD). pp. 220–227 (1996)
53. Smolensky, R.: Algebraic methods in the theory of lower bounds for boolean circuit complexity. In: Proc. 19th ACM Symposium on Theory of Computing. pp. 77–82. ACM Press (1987)
54. Staber, S., Bloem, R.: Fault localization and correction with QBF. In: International Conference on Theory and Applications of Satisfiability Testing (SAT) 2007. pp. 355–368 (2007)
55. Stockmeyer, L.J., Meyer, A.R.: Word problems requiring exponential time: Preliminary report. In: Annual Symposium on Theory of Computing. pp. 1–9. ACM (1973)
56. Vardi, M.Y.: Boolean satisfiability: Theory and engineering. Commun. ACM 57(3), 5 (2014)
57. de la Vega, W.F.: On random 2-SAT (revised version) (1998), preprint
58. de la Vega, W.F.: Random 2-SAT: results and problems. Theoretical Computer Science 265(1-2), 131–146 (2001)
59. Vollmer, H.: Introduction to Circuit Complexity – A Uniform Approach. Texts in Theoretical Computer Science, Springer Verlag, Berlin Heidelberg (1999)

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il