Electronic Colloquium on Computational Complexity, Report No. 219 (2026)

Fine-Grained Size-Cost-Capacity for Strong Semantic QBF Proof Systems

Olaf Beyersdorff
Friedrich Schiller University Jena, Germany

Lea Kasche
Friedrich Schiller University Jena, Germany

Luc Nicolas Spachmann
Friedrich Schiller University Jena, Germany

Abstract

We revisit the semantic size-cost-capacity technique (Beyersdorff, Blinkhorn & Hinde, 2019) for proof-size lower bounds in proof systems for quantified Boolean formulas (QBF). While the original technique is only applicable to weak proof systems with bounded capacity, we present a fine-grained generalisation of this technique that allows us to attack stronger systems with unbounded capacity. Both the technique and the relevant QBF proof systems are defined semantically and are parametrised by the proof lines permitted in the system.

We demonstrate our approach on the Q-k-DNF systems — allowing for k-DNFs as lines and forming a QBF analogue of the propositional Res(k) proof systems (Krajíček, 2001) — and show that the Q-k-DNF hierarchy is strict with exponential separations between all levels, using our refined size-cost-capacity technique.

2012 ACM Subject Classification Theory of computation → Proof complexity

Keywords and phrases QBF, proof complexity, Res(k), lower bounds

Funding *Olaf Beyersdorff*: Carl-Zeiss Foundation and DFG grant BE 4209/3-2
Lea Kasche: Carl-Zeiss Foundation

1 Introduction

Proof complexity studies the resources needed to prove theorems in formal calculi, with proof size as the primary measure. The field bears tight connections to computational complexity [15], to mathematical logic via bounded arithmetic [14, 22], and to solving: each solver run on an unsatisfiable instance can be interpreted as a proof, whereby proof-size lower bounds translate into lower bounds on solver running time [12, 30].

The success story of SAT solving [24] has been lifted to the computationally even harder setting of quantified Boolean formulas (QBF). As the prototypical PSPACE-complete problem, QBFs succinctly encode problems from verification, planning, and synthesis [28]. Conceptually different solving paradigms coexist: QCDCL [17, 31] and CEGAR approaches as in RAReQS [20] and CAQE [25], mirrored proof-theoretically by a number of QBF calculi [3, 19, 21]. The relative strength of these QBF resolution systems is well understood [6]; notably, their lower bounds require techniques quite different from propositional methods (cf. [2, 9] for overviews).

Many of the strongest QBF proof systems arise uniformly by augmenting a propositional base system P with the universal reduction rule, yielding $P + \forall\text{red}$ [7] — a spectrum reaching from QU-Res [29] up to QBF Frege systems. Dedicated lower-bound techniques for these systems include strategy extraction, turning circuit lower bounds into proof-size bounds [7]; the semantic size-cost-capacity technique [5], recently employed to construct hard QBFs from computationally hard problems [26]; and the characterisation of QU-Res hardness through

ISSN 1433-8092

unified decision lists [10]. All aim for 'genuine' QBF lower bounds not simply inherited from the base system [8, 13].

Of these techniques, size-cost-capacity [5] is arguably the most elegantly semantic: every $P + \forall\text{red}$ refutation π of a false QBF Φ satisfies $|\pi| \geq \text{cost}(\Phi)/\text{capacity}(\pi)$, where cost counts the responses each universal winning strategy must provide for some universal block, and capacity bounds the responses a single proof line can contribute. The technique is highly effective when capacity is bounded — as for QU-Res and CP + $\forall$red, where capacity is 1 [5] — but therein lies its fundamental limitation: in stronger systems capacity explodes. This holds for QBF Frege, where the technique provably fails [5], but already for QBF versions of Res(k) [23] with k-DNFs as lines. Thus size-cost-capacity is silent on almost the entire landscape between QBF resolution and QBF Frege.

Our Contributions. We develop a fine-grained version of the size-cost-capacity technique that yields lower bounds for strong semantic QBF proof systems on which the original method fails.

1. Semantic QBF proof systems. We take a very general view on QBF systems within the reduction framework: starting from a syntactic representation class P for lines, we directly define the semantic proof system $\mathcal{Q}$-P, combining semantic entailment and universal reduction (Section 3). These systems are sound and complete by design, and they subsume the semantic variants of the known reduction-based calculi; e.g. semantic QU-Res coincides with $\mathcal{Q}$-clauses. Lower bounds for $\mathcal{Q}$-P thus apply to every calculus operating with lines from P, irrespective of its rules, and are genuine in the sense of Beyersdorff et al. [8].

2. Fine-grained size-cost-capacity. We generalise the size-cost-capacity paradigm by parametrising it with a function representation class $\mathcal{L}$: response maps no longer return concrete assignments, but functions from $\mathcal{L}$ that compute the responses, and winning strategies are replaced by *Two-Step-Strategies*, from which we extract responses in two steps. Our main result (Theorem 3.10) shows that $|\pi| \geq \text{cost}_{\mathcal{L}}(\Phi)/\text{capacity}_{\mathcal{L}}(\pi)$ for every $\mathcal{Q}$-P refutation π of Φ. Intuitively, the parametrisation rescales cost and capacity simultaneously: choosing $\mathcal{L}$ suitably brings both measures into a regime where capacity becomes polynomial while cost remains exponential. This makes the technique applicable in scenarios where capacity in the sense of Beyersdorff et al. [5] is unbounded; for $\mathcal{L}$ the class of constant functions we exactly recover the original technique.

3. The $\mathcal{Q}$-k-DNF hierarchy is strict. As a test case we apply our framework to $\mathcal{Q}$-k-DNF, the semantic QBF analogues of Res(k) (Section 4). Here the original capacity is still exponential, but the rescaled k-CNF-capacity equals 1. We introduce the QBFs k-`Equality`, generalising the equality formulas of Beyersdorff et al. [5], and show via a new parity-covering lemma that their k-CNF-cost is exponential. Together this yields exponential lower bounds for k-`Equality` in $\mathcal{Q}$-k-DNF, whereas the formulas admit linear-size $\mathcal{Q}$-$(k+1)$-DNF refutations. Hence consecutive levels of the $\mathcal{Q}$-k-DNF hierarchy are exponentially separated — a quantified counterpart to the strictness of the propositional Res(k) hierarchy [16, 27], but for semantic systems, where propositional techniques are unavailable.

4. Beyond k-DNFs. Our approach in principle applies to further strong QBF systems, e.g. semi-algebraic QBF systems [11]. As a second example, we consider as lines degree-k polynomial threshold functions (PTF$_k$), yielding a hierarchy of systems $\mathcal{Q}$-PTF$_k$ above semantic QBF cutting planes [4] that p-simulate the corresponding levels $\mathcal{Q}$-k-DNF (Section 5).

Organisation. We review preliminaries in Section 2. Section 3 introduces the semantic $\mathcal{Q}$-P framework and establishes the fine-grained size-cost-capacity theorem. Section 4 applies

the technique to the $\mathcal{Q}$-k-DNF hierarchy, and Section 5 to threshold systems. We conclude in Section 6 with a discussion on further directions.

2 Preliminaries

For $n \in \mathbb{N}$, we use $[n]$ to denote the set $\{1, \dots, n\}$. We assume familiarity with basic notions from computational complexity, cf. [1], as well as from logic, cf. [22], but define all specific concepts needed in this paper.

We consider propositional formulas φ built from constants $0, 1$, connectors $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$, and propositional variables. A literal is a variable v or its negation $\neg v$. A clause is a disjunction of literals, and a formula is in Conjunctive Normal Form (CNF) if it is a conjunction of clauses. Dually, a term is a conjunction of literals, and a formula is in Disjunctive Normal Form (DNF) if it is a disjunction of terms. For some natural number k, we say a formula is in k-DNF (resp. k-CNF), if it is in DNF (resp. CNF) and every term (resp. clause) consists of at most k literals. When V is a set of variables, a (partial) assignment to V is a (partial) function $\alpha : V \to \{0, 1\}$. We write $\langle V \rangle$ for the set of all complete assignments to V, and $\mathrm{vars}(\varphi)$ or $\mathrm{vars}(\alpha)$ for the set of all variables occurring in φ or α (i.e. variables assigned truth values by α). If $V = \{x_1, \dots, x_n\}$ we also write $\langle x_1, \dots, x_n \rangle$ for $\langle V \rangle$. We denote by $\varphi|_\alpha$ the formula φ where each variable $v \in \mathrm{vars}(\alpha)$ has been substituted by $\alpha(v)$. When a function is applied to an assignment of a superset of its variables, we implicitly project the assignment to the function's domain.

Quantified Boolean Formulas. A (closed prenex) *Quantified Boolean Formula* (QBF) is a formula $Q.\varphi$ where φ is a propositional formula in CNF and Q is the quantifier prefix that quantifies all variables v in φ either existentially as $\exists v$ or universally as $\forall v$. Whether a QBF is true or not can be defined recursively. The formula $\forall u\, Q.\varphi$ is true if both $Q.\varphi|_{u=0}$ and $Q.\varphi|_{u=1}$ are true. The formula $\exists x\, Q.\varphi$ is true if at least one of $Q.\varphi|_{x=0}$ and $Q.\varphi|_{x=1}$ is true.

In a fully quantified prenex QBF, the quantifier prefix determines a total order of the variables. We will sometimes refer to the variables preceding a variable v in the prefix as variables *left* of v; analogously for the variables *right* of v.

A QBF $Q_1 x_1 \dots Q_k x_k.\varphi$ can be seen as a game between two players: *universal* ($\forall$) and *existential* ($\exists$). In the i-th step of the game, the player Q_i assigns a value to the variable x_i. The existential player wins if φ evaluates to 1 under the constructed assignment, otherwise the universal player wins. Given a universal variable u with index i, a *strategy for* u is a function that maps each assignment of the existential variables of index $< i$ to a value in $\{0, 1\}$. A QBF is false if and only if there exists a *winning strategy* for the universal player, that is if the universal player has a strategy for all universal variables that wins any possible game [1, 18]. Formally, this universal winning strategy is a set of functions f_i, one for every universal variable u_i, from the assignments of all previous existential variables to $\{0, 1\}$. We require that for every assignment α to the existential variables, replacing all u_i by $f_i(\alpha)$ and x_i by $\alpha(x_i)$ must falsify the matrix.

Proof systems. According to Cook and Reckhow [15], a *proof system* for a language $\mathcal{L}$ is a polynomial-time onto function $P : \{0, 1\}^* \to \mathcal{L}$. Each string $\varphi \in \mathcal{L}$ is a *theorem* and if $P(\pi) = \varphi$, then π is a *proof* of φ in P (a P-proof). Given a polynomial-time function $P : \{0, 1\}^* \to \{0, 1\}^*$ the fact that $P(\{0, 1\}^*) \subseteq \mathcal{L}$ is the *soundness* of P and the fact that $P(\{0, 1\}^*) \supseteq \mathcal{L}$ is the *completeness* of P. A proof system P_1 *p-simulates* a different proof system P_2, if every P_2 proof can be transformed into a P_1 proof of the same theorem in polynomial-time.

3 The refined lower bound technique

3.1 The $\mathcal{Q}$-P proof system framework

We start by restating and slightly expanding the general framework of reduction-based QBF proof systems of Beyersdorff et al. [5]. To achieve this, we first need to define a suitable Boolean function representation class, called *base representation class*.

▶ **Definition 3.1** (Base representation class). *Let P be some Boolean function representation class. We say P is a* base representation class*, if it fulfils the following conditions:*

1. *Every clause C with k literals can be efficiently encoded into a line of P, i.e. $\mathrm{enc}_P(C) \in P$ with $|\mathrm{enc}_P(C)| \in \mathrm{poly}(k)$,*
2. *for every $f \in P$ and partial assignment α to the variables of f, $f|_\alpha$ can be computed efficiently.*

Less formally, condition 1 requires these representations to be able to efficiently represent clauses, and condition 2 enables conditioning.

▶ **Definition 3.2** ($\mathcal{Q}$-P). *Let P be some base function representation class and Φ a false closed prenex QBF. We define a $\mathcal{Q}$-P proof π of Φ to be a sequence of lines $\ell_1, \ldots, \ell_s$ with $\ell_i \in P$ for $i \in [s]$, such that $\ell_s \equiv \bot$ and for every $i \in [s]$:*

- *$\ell_i = \mathrm{enc}_P(C)$ for a clause C of φ (axiom), or*
- *there exists $1 \leq i_1, \ldots, i_t < i$ such that $\ell_{i_1}, \ldots, \ell_{i_t} \models \ell_i$ (semantic entailment), or*
- *there exists $1 \leq j < i$, a universal variable u occurring in the rightmost quantifier block among the variables of ℓ_j, and a Boolean value b, such that $\ell_j|_{u=b} = \ell_i$ (universal reduction).*

The size of π is defined by $\mathrm{size}(\pi) = |\pi| = \sum_{i \in [s]} |\ell_i|$, where $|\ell_i| \geq 1$ denotes the representation size of ℓ_i; we also write $|\pi|$ for $\mathrm{size}(\pi)$.

▶ **Proposition 3.3** ($\mathcal{Q}$-P is sound and complete). *Let P be a base representation class. Then a QBF has a $\mathcal{Q}$-P refutation if and only if it is false.*

Proof. *Completeness.* Let Φ be a false QBF and let $\pi = (C_1, \ldots, C_s)$ be a semantic QU-Res refutation of Φ; such a refutation exists by the completeness of the semantic framework of Beyersdorff et al. [5]. We translate π line by line: if C_i is an axiom or derived by semantic entailment, let $D_i = \mathrm{enc}_P(C_i)$, which exists by condition 1 of Definition 3.1; if $C_i = C_j|_{u=b}$ was derived by universal reduction, let $D_i = D_j|_{u=b}$, which lies in P by condition 2. In all cases $D_i \equiv C_i$, so each step of π' is a valid application of the corresponding $\mathcal{Q}$-P rule and $\pi' = (D_1, \ldots, D_s)$ is a $\mathcal{Q}$-P refutation of Φ.

Soundness. Let now $\Phi = Q.\varphi$ be a true QBF, S an existential winning strategy and assume for the sake of contradiction that there is a $\mathcal{Q}$-P refutation $\pi = (D_1, \ldots, D_s)$ of Φ. Since S is an existential winning strategy, $(\alpha, S(\alpha)) \models \varphi$ for all universal assignments α. We show inductively over i that $(\alpha, S(\alpha)) \models D_i$ for all α.

- If D_i is an axiom, $(\alpha, S(\alpha)) \models D_i$ because $D_i \equiv C$ for a clause C of φ.
- If D_i was derived using semantic entailment from $D_{i_1}, \ldots, D_{i_t}$, then every assignment satisfying $D_{i_1}, \ldots, D_{i_t}$ also satisfies D_i. The induction hypothesis yields $(\alpha, S(\alpha)) \models D_i$.
- If $D_i = D_j|_{u=b}$ was derived using universal reduction from D_j on variable u, let α' be α with the value of u flipped. All existential variables of D_j are left of u, so $S(\alpha)$ and $S(\alpha')$ agree on the existential variables of D_j; hence $(\alpha, S(\alpha))$ and $(\alpha', S(\alpha'))$ agree on $\mathrm{vars}(D_j) \setminus \{u\}$. By the induction hypothesis, $(\alpha, S(\alpha)) \models D_j$ and $(\alpha', S(\alpha')) \models D_j$. One of the two assignments sets $u = b$, and therefore $(\alpha, S(\alpha)) \models D_j|_{u=b} = D_i$.

In particular, $(\alpha, S(\alpha)) \models D_s \equiv \bot$, a contradiction. ◀

The original definition from Beyersdorff et al. [5] required P to be a suitable propositional proof system, called base system, which was then augmented by the universal reduction rule. Subsequently, there has been a lot of interest in *genuine* QBF lower bounds, studying proof systems with the semantic entailment and universal reduction rules and using the lines of the propositional base system. However, in these semantic QBF proof systems, the rules of the propositional base systems are immaterial and we can restrict the definition to the type of lines used, resulting in our definition above. As an example, semantic QU-Res is equivalent to $\mathcal{Q}$-clauses.

3.2 Defining $\mathcal{L}$-size-cost-capacity

In this section, we will introduce a generalisation of the size-cost-capacity framework described in Beyersdorff et al. [5]. The original framework can be used to prove lower bounds for a range of proof systems that have low capacity. However, many other proof systems have exponential capacity, which usually precludes any interesting results from this technique. We will parametrize the definitions of cost and capacity by a function representation class $\mathcal{L}$. Picking a more expressive class $\mathcal{L}$ will decrease both size and cost.

Interesting lower bounds for a formula and proof system can be obtained by choosing a class $\mathcal{L}$ of just the right strength such that a proof system's capacity is small (e.g. constant or polynomial), while the cost of the formula barely stays large (e.g. exponential). This requires a careful trade-off when it comes to the expressivity of the formula class.

We start by introducing the necessary definitions required for our main result. These definitions are generalisations of the definitions in Beyersdorff et al. [5]. Following Beyersdorff et al. [5], the *rightmost block* of a line L with respect to a quantifier prefix $\mathcal{Q}$ is the rightmost quantifier block of $\mathcal{Q}$ containing a variable of L; we call L *reducible* if its rightmost block is universal.

▶ **Definition 3.4** ($\mathcal{L}$-Response Map). *Let P be a base representation class, L be a reducible line in a $\mathcal{Q}$-P refutation of a QBF $Q.\varphi$ with rightmost (universal) block U, $X = \mathrm{vars}(L) \setminus U$ and $\mathcal{L}$ some (possibly nondeterministic) function representation class. An $\mathcal{L}$-Response Map $\mathcal{R}$ of L with respect to $\mathcal{Q}$ is a function*

$$\mathcal{R} : \langle X \rangle \to \mathcal{L},$$

where for each $\alpha \in \langle X \rangle$, $\mathcal{R}(\alpha)$ represents a (possibly partial, nondeterministic) function into $\langle U \rangle$ satisfying the following condition:

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

Note that if $\mathcal{L}$ represents a nondeterministic class, the condition has to hold for every possible output.

The intuition here is that $\mathcal{R}(\alpha)$ is not the actual response on the input α, but a function that can compute this response given α. This enables us to group several inputs together if their responses can be computed by a single function from $\mathcal{L}$, reducing the range of $\mathcal{R}$. The more expressive the class $\mathcal{L}$ is, the more responses can be grouped together this way. If $\mathcal{L}$ can represent arbitrary functions, then it is always possible to have one function representing every response at once, so $|\mathrm{rng}(\mathcal{R})|$ can be 1. We are therefore interested in less expressive representation classes $\mathcal{L}$.

Combining $\mathcal{L}$-response maps for every reducible line in a proof yields a *response map set.*

▶ **Definition 3.5** ($\mathcal{L}$-Response Map Set). *Given a $\mathcal{Q}$-P refutation π of a QBF $Q.\varphi$ whose reducible lines are $L_1, \ldots, L_r$, an $\mathcal{L}$-response map set for π is a set $\{\mathcal{R}_1, \ldots, \mathcal{R}_r\}$ in which each $\mathcal{R}_i$ is an $\mathcal{L}$-response map for L_i with respect to Q.*

The $\mathcal{L}$-capacity is a measure of complexity for a proof and proof system. Informally, it measures how hard it is to falsify a single line in the proof.

▶ **Definition 3.6** ($\mathcal{L}$-Capacity). *Let π be a $\mathcal{Q}$-P refutation of a QBF $Q.\varphi$. The $\mathcal{L}$-capacity of an $\mathcal{L}$-response map set $\mathcal{R} = \{\mathcal{R}_1, \ldots, \mathcal{R}_r\}$ for π is*

$$\text{capacity}_{\mathcal{L}}(\mathcal{R}, \pi) = \max_{i \in [r]} \Big\{ \sum_{L \in \text{rng}(\mathcal{R}_i)} |L| \Big\}.$$

The $\mathcal{L}$-capacity of π is the minimal capacity of any $\mathcal{L}$-response map set for π.

Next, we introduce the Two-Step-Strategy. In essence, it is a way to write down a universal winning strategy, using some function representation class.

▶ **Definition 3.7** (Two-Step-Strategy). *Let $\Phi = \exists X_1 \forall U_1 \exists X_2 \ldots \exists X_d \forall U_d \exists X_{d+1}.\varphi$ be a false QBF and $\mathcal{L}$ some function representation class. A* Two-Step-Strategy *T for Φ and $\mathcal{L}$ is a tuple of functions $(T_1, \ldots, T_d)$, such that*

- *$T_i : \langle X_1 \cup \cdots \cup X_i \rangle \to \mathcal{L}$, with $T_i(\alpha) : \langle X_1 \cup U_1 \cup \cdots \cup U_{i-1} \cup X_i \rangle \to \langle U_i \rangle$,*
- *every possible value of $T_i(\alpha)(\alpha, \beta_i(\alpha))$ is defined, where $\beta_i : \langle X_1 \cup \cdots \cup X_{i-1} \rangle \to \langle U_1 \cup \cdots \cup U_{i-1} \rangle$ are the responses of $T_1, \ldots, T_{i-1}$, and for every choice of such values the resulting functions $S_i(\alpha) \in \langle U_i \rangle$ form a universal winning strategy $S = \bigcup_i S_i$ for Φ.*

We define $\text{size}(T_i) = \sum_{L \in \text{rng}(T_i)} |L|$ *and the cost* $\text{cost}(T) = \max_i \text{size}(T_i)$.

The $\mathcal{L}$-cost of a QBF measures the hardness of a Two-Step-Strategy for a single universal block.

▶ **Definition 3.8** ($\mathcal{L}$-Cost). *Let Φ be a false QBF and $\mathcal{L}$ some function representation class. The $\mathcal{L}$-Cost of Φ is defined as*

$$\text{cost}_{\mathcal{L}}(\Phi) = \min_{T} \text{cost}(T),$$

where the minimum ranges over all Two-Step-Strategies T for Φ and $\mathcal{L}$.

To illustrate that these definitions are indeed generalisations of Beyersdorff et al. [5], let us compare them. In their paper, they did not include the intermediate step of function representations (our parameter $\mathcal{L}$), such that response maps directly map to assignments. We can simulate this behaviour by setting $\mathcal{L}$ to be constant functions, denoted by c. Additionally, we can reconstruct response maps and universal winning strategies from $\mathcal{L}$ response maps and Two-Step-Strategy respectively by evaluating the expressed function on the input: Given any $\mathcal{L}$-response map $\mathcal{R}$ for a line L in a $\mathcal{Q}$-P proof π, defining

$$O(\alpha) := \mathcal{R}(\alpha)(\alpha)$$

for all α is a (non-parametrised) response map for L and π. We call this the *underlying response map*. Analogously, we define the *underlying response map set*, and the *underlying strategy*. Note that the range of underlying response maps and strategies and, as such, the capacity and cost, can increase exponentially compared to the parametrised versions.

With these definitions, they stated the size-cost-capacity theorem, which is identical to Theorem 3.10 with $\mathcal{L}$ set to constants. We will use the same techniques to obtain this

result in the parametrised variant. The more succinct representations using the function representation classes allow us to study proof systems with exponential capacity, where their technique fails. To prove the theorem, we adapt the round-based strategy extraction to extract a Two-Step-Strategy.

▶ **Theorem 3.9** (Round-based Two-Step-Strategy extraction). *Let π be a $\mathcal{Q}$-P refutation of a QBF $\Phi = \exists X_1 \forall U_1 \exists X_2 \dots \exists X_d \forall U_d \exists X_{d+1}.\varphi$ with reducible lines $L_1, \dots, L_r$ and a $\mathcal{L}$-response map set $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_r\}$. Further, for each $\alpha \in \langle \text{vars}_\exists(\Phi) \rangle$ and $i \in [d]$:*

- *let σ_0^α be the empty assignment;*
- *let α_i be the projection of α to X_i;*
- *let F_i^α be the first line in π, such that $\text{vars}(F_i^\alpha) \subseteq \bigcup_{j \le i}(X_j \cup U_j)$ and $F_i^\alpha|_{\sigma_{i-1}^\alpha \cup \alpha_i}$ is non-tautological;*
- *if F_i^α is some $L_{j_i^\alpha}$ whose rightmost block is U_i, then $\beta_i^\alpha = \mathcal{R}_{j_i^\alpha}(\sigma_{i-1}^\alpha \cup \alpha_i)$, otherwise let $\beta_i^\alpha = \mathcal{R}_{j_i^U}(\sigma_{i-1}^\alpha \cup \alpha_i)$, where j_i^U is the minimal index for which U_i is the rightmost block of $L_{j_i^U}$;*
- *let $\sigma_i^\alpha = \sigma_{i-1}^\alpha \cup \alpha_i \cup (\beta_i^\alpha)(\sigma_{i-1}^\alpha \cup \alpha_i)$.*

The extracted Two-Step-Strategy *for π with respect to $\mathcal{R}$ is $T = (T_i)_{i \in [d]}$ with*

$$T_i : \langle X_1 \cup \dots \cup X_i \rangle \to \mathcal{L}$$
$$\bigcup_{j \in [i]} \alpha_j \mapsto \beta_i^\alpha.$$

This is a Two-Step-Strategy for Φ.

Proof. The only difference to the existing strategy extraction framework of Beyersdorff et al. [5] is that we are using an $\mathcal{L}$-response map instead of a standard response map. As a consequence, the proof is analogous.

For every i and α, $\text{vars}(\bot) \subseteq \bigcup_{j \in [i]}(X_j \cup U_j)$ and $\bot|_{\sigma_{i-1}^\alpha \cup \alpha_i}$ is non-tautological. As such, every F_i^α exists. We can assume without loss of generality, that every universal variable appears in some axiom; this yields the existence of j_i^U.

Every $\mathcal{L}$-response $\mathcal{R}(\alpha)$ can be turned into a standard response $(\mathcal{R}(\alpha))(\alpha)$. In the original formulation, β_i^α was a direct assignment to the universal variables of the ith block; now it is a function that takes an assignment to the previous variables, and $(\beta_i^\alpha)(\sigma_{i-1}^\alpha \cup \alpha_i)$ yields the assignment.

Using this translation, we can use Beyersdorff et al. [5, Lemma 5.11] that the underlying strategy of our extracted Two-Step-Strategy is correct. This is sufficient for the correctness of a Two-Step-Strategy. (Note that Proposition 5.9 and Proposition 5.10 in Beyersdorff et al. [5] hold verbatim for our slightly changed $\mathcal{Q}$-P framework).

◀

▶ **Theorem 3.10** ($\mathcal{L}$-Size-Cost-Capacity). *Let π be a $\mathcal{Q}$-P refutation of a QBF Φ. Then*

$$|\pi| \ge \frac{\text{cost}_\mathcal{L}(\Phi)}{\text{capacity}_\mathcal{L}(\pi)}.$$

Proof. Let $\Phi = \exists X_1 \forall U_1 \dots \exists X_d \forall U_d \exists X_{d+1}.\varphi$ (w.l.o.g. of this form, padding with empty blocks if necessary). Let $L_1, \dots, L_r$ be the reducible lines of π and let $\mathcal{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_r\}$ be an $\mathcal{L}$-response map set for π of minimum capacity. Further, let T be the extracted Two-Step-Strategy for π with respect to $\mathcal{R}$. By Theorem 3.9, T is a valid Two-Step-Strategy for Φ. Hence there is some universal block U_i such that

$$\text{size}(T_i) \ge \text{cost}_\mathcal{L}(\Phi).$$

By construction, every function in $\mathrm{rng}(T_i)$ is chosen from the range of a response map $\mathcal{R}_j$ attached to a reducible line L_j of π. There are at most $|\pi|$ reducible lines, and for each of them $\sum_{L\in\mathrm{rng}(\mathcal{R}_j)}|L| \leq \mathrm{capacity}_{\mathcal{L}}(\pi)$. Thus

$$\mathrm{cost}_{\mathcal{L}}(\Phi) \leq \mathrm{size}(T_i) = \sum_{L\in\mathrm{rng}(T_i)} |L| \leq |\pi| \cdot \mathrm{capacity}_{\mathcal{L}}(\pi),$$

which is the claimed bound. ◀

4 Applying our technique to $\mathcal{Q}$-k-DNF

We apply this framework to k-DNF as lines for $k \geq 1$.

▶ **Lemma 4.1.** *For all $k \geq 1$, k-DNF are a base representation class.*

Proof. Clauses are 1-DNFs. As such, for condition 1, $\mathrm{enc}(C) = C$. Restricting k-DNF yield k-DNF, as such Condition 2 follows directly. ◀

We can then use Definition 3.2 to define the proof system $\mathcal{Q}$-k-DNF. To recapitulate, this proof system uses k-DNF as lines, can download clauses as axioms and has semantic entailment and universal reduction as derivation rules.

If Φ has n variables, every k-DNF over these variables consists of at most $\sum_{j=0}^{k} 2^j \binom{n}{j} \leq \mathrm{poly}(n)$ distinct terms. As we are interested in exponential separations, these polynomial factors are not relevant and we can define the size of every k-DNF (and analogously k-CNF) to be 1.

$\mathcal{Q}$-k-DNF also show the advantage of our new technique compared to the original size-cost-capacity theorem. For $k \geq 2$, the capacity of $\mathcal{Q}$-k-DNF can be exponential. As such, it could not be used to show any lower bounds and separations within the hierarchy. With our technique, we will show an exponential separation between $\mathcal{Q}$-k-DNF and $\mathcal{Q}$-$(k+1)$-DNF. To achieve this, we extend the well-known class of QBF `Equality` [5] to ℓ-`Equality` as the families of separating formulas.

▶ **Definition 4.2.** *Let $X = \{x_i^j \mid i \in [n], j \in [\ell]\}$, $U = \{u_i \mid i \in [n]\}$ and $T = \{t_i \mid i \in [n]\}$ be disjoint sets of Boolean variables. Then, the ℓ-$\mathtt{Equality}_n$ formula is defined as*

$$\exists X \forall U \exists T.$$
$$\bigwedge_{i\in[n]} \left(((\bigoplus_{j\in[\ell]} x_i^j) \leftrightarrow u_i) \rightarrow t_i \right) \wedge \left(\bigvee_{i\in[n]} \neg t_i \right).$$

This can be written as a CNF with a polynomial blow-up in n and without additional variables. It has an exponential blow-up in ℓ, but ℓ is constant. Note that 1-`Equality` is equivalent to `Equality`. This formula has a unique winning strategy for the universal player, which sets every u_i to the parity of the x_i^j over all j.

We will first argue that k-`Equality` is easy for $\mathcal{Q}$-$(k+1)$-DNF. The main idea here is that the semantics of k-`Equality` can be expressed in a single $(k+1)$-DNF. The semantics we care about here are obtained by existentially quantifying T in the matrix of k-`Equality`, leaving a formula in X and U which is equivalent to the formula L below. Expressing these semantics in a single line enables using universal reduction to subsequently eliminate all variables from U.

On the other hand, expressing the semantics of k-`Equality` using a single k-DNF is impossible. Moreover, even polynomially many lines are not enough; capturing the semantics

(in a more informal sense) requires exponentially many k-DNF lines. We will formalize both of these arguments next.

▶ **Lemma 4.3.** *For all $k \geq 1$, $\mathcal{Q}$-$(k+1)$-DNF have short refutations for k-*`Equality`*.*

Proof. The formula

$$L := \bigvee_{i\in[n]} \left(\left(\bigoplus_{j\in[k]} x_i^j \right) \not\leftrightarrow u_i \right)$$

can be written as single $(k+1)$-DNF with $2^k n$ terms. Simultaneously, L is semantically implied by the matrix of k-`Equality` and contains no variables of the last quantifier block. As such, we can perform universal reductions step by step on all universal variables. We define

$$L_m = \bigvee_{i\in[m]} \left(\left(\bigoplus_{j\in[k]} x_i^j \right) \not\leftrightarrow u_i \right)$$

for $m \in \{0, \dots, n\}$. Obviously, $L_n = L$ and $L_0 = \bot$. Additionally, for all $m \in [n]$,

$$L_m|_{u_m=0} = L_{m-1} \vee \left(\bigoplus_{j\in[k]} x_m^j \leftrightarrow 1 \right),$$
$$L_m|_{u_m=1} = L_{m-1} \vee \left(\bigoplus_{j\in[k]} x_m^j \leftrightarrow 0 \right) \text{ and}$$
$$(L_m|_{u_m=0}, L_m|_{u_m=1}) \models L_{m-1}.$$

As such, let $C_1, \dots, C_r$ be the clauses of a CNF encoding of the matrix of k-`Equality`. Since k is fixed, $r = O(n)$, and the clauses $C_1, \dots, C_r$ can be introduced as axioms. The line L_n is then derived from these axioms by semantic entailment. Thus

$$\pi = (C_1, \dots, C_r, L_n, L_n|_{u_n=0}, L_n|_{u_n=1}, L_{n-1}, \dots, L_1, L_1|_{u_1=0}, L_1|_{u_1=1}, L_0 = \bot)$$

is a $\mathcal{Q}$-$(k+1)$-DNF refutation of k-`Equality` of size $\Theta(n)$. ◀

To show the hardness result for $\mathcal{Q}$-k-DNF, we first need to introduce a suitable function representation class.

▶ **Definition 4.4.** *Let X, U be disjoint sets of Boolean variables and let $f : \langle X \cup U \rangle \to \{0,1\}$ be (a representation of) a Boolean function. We define the* computation variant *of f to be the possibly multi-valued partial function*

$$f^{\text{comp}} : \langle X \rangle \to \langle U \rangle,$$
$$f^{\text{comp}}(x) \in \{u \in \langle U \rangle \mid f(x,u) = 1\}.$$

In general, this function is nondeterministic. If $f|_x$ is satisfiable, every satisfying assignment $u \in \langle U \rangle$ is a possible output.

We will use the computation variant of k-CNF as our representation class. For readability, we will omit the superscript.

We start with the upper bound of the capacity.

▶ **Lemma 4.5.** *Every $\mathcal{Q}$-k-DNF proof π has* $\text{capacity}_{k\text{-}CNF}(\pi) = 1$.

Proof. Every line $L \in \pi$ is per definition a k-DNF. As such, the negation of L is a k-CNF $\overline{L}$. Note that we regard $\overline{L}$ as a function into the whole rightmost universal block, not restricted to the variables appearing in L. In particular, $\mathcal{R} : \alpha \mapsto \overline{L}$ is a k-CNF-Response Map for L. This holds for every line of π and $\text{capacity}_{k\text{-CNF}}(\pi) = 1$. ◀

To show the cost lower bound, we use the following coding lemma.

▶ **Lemma 4.6** (Parity-covering lemma). *Let $X = \{x_i^j \mid i \in [n], j \in [k]\}$, $U = \{u_i \mid i \in [n]\}$ be the variables of k-$\texttt{Equality}_n$, and let $p : \langle X \rangle \to \langle U \rangle$ be the unique universal winning strategy*

$$p(\alpha) = (u_1, \ldots, u_n), \text{ with}$$
$$u_i = \bigoplus_{j \in [k]} \alpha(x_i^j).$$

Consider a k-CNF $F(X, U)$ and let $B = B_F$ be the set of existential assignments to which it is a valid reply, i. e.

$$B_F = \big\{\alpha \in \langle X \rangle \mid F(\alpha, p(\alpha)) = 1 \text{ and } F(\alpha, \beta) = 0 \text{ for all } \beta \neq p(\alpha)\big\}.$$

Then $|B| \leq 2^{kn - n/k^2}$.

Proof. We refer to the variables $\{x_i^1, \ldots, x_i^k\}$ as *block* i, and also associate the variable u_i with block i. Consider an $\alpha \in B$ and $i \in [n]$ and write $p(\alpha)^{\oplus i}$ for the assignment obtained by flipping the value of u_i in $p(\alpha)$. We know that $F(\alpha, p(\alpha)^{\oplus i}) = 0$ due to the definition of B, so there is a clause $C_{\alpha,i} \in F$ that is falsified by $(\alpha, p(\alpha)^{\oplus i})$. We choose an arbitrary such $C_{\alpha,i}$ and call it the critical clause for α and block i. Because $F(\alpha, p(\alpha)) = 1$, the critical clause is satisfied by $(\alpha, p(\alpha))$, so it has to contain u_i in the same polarity as $p(\alpha)$ and all its other literals are falsified by $(\alpha, p(\alpha))$. Because F is a k-CNF, $C_{\alpha,i}$ contains at most $k - 1$ variables other than u_i, and therefore omits at least one of the variables $\{x_i^1, \ldots, x_i^k\}$.

We now construct a randomized assignment γ depending on F. We start by uniformly choosing a permutation σ of the blocks $[n]$, and independently for each block i, a variable $z_i \in \{x_i^1, \ldots, x_i^k\}$. We iteratively build the assignment γ, processing the blocks in the order σ. In each block, we guess the $k - 1$ variables that are not z_i uniformly at random, and then test both possible values of z_i. A value b is called *locally refuted* if some clause of F has all its literals falsified by the assignment constructed so far extended by $z_i = b$, where each u_m for an already processed block m (including $m = i$) takes the value $\bigoplus_{j \in [k]} \gamma(x_m^j)$ of the current assignment (using $z_i = b$ for block i). If exactly one value is locally refuted, we set z_i to the other value, otherwise we pick a uniform random value. We do this for each block in order to obtain γ.

Consider a fixed block i, $\alpha \in B$, and randomisation outcome. Let $C_{\alpha,i}$ be the critical clause. We say that block i is *good* for α if both of the following conditions hold:

1. among the blocks with variables in $C_{\alpha,i}$, block i is last according to σ.
2. z_i does not occur in $C_{\alpha,i}$.

Because $C_{\alpha,i}$ has variables from at most k blocks, the first condition is satisfied with probability at least $1/k$ (depending on the choice of σ). Because the clause omits at least one of the variables $\{x_i^1, \ldots, x_i^k\}$, the second condition is satisfied with probability at least $1/k$ (depending on the choice of z_i). Because these choices are independent, the probability that block i is good for α is at least $1/k^2$.

We now argue that after guessing all earlier variables according to α, if block i is good for α, we will decide z_i according to α without a random choice. All literals of $C_{\alpha,i}$ except the

one on u_i are already determined according to α, because they are decided before z_i (and z_i itself does not occur in the clause). In particular, they are falsified by $(\alpha, p(\alpha))$. The true value $b = \alpha(z_i)$ cannot be locally refuted, because $(\alpha, p(\alpha))$ satisfies F. In the clause $C_{\alpha,i}$, only the literal on u_i is satisfied. Flipping the value of z_i flips the value of u_i and does not change any other literal in that clause, so it will be falsified and $b = \neg\alpha(z_i)$ is locally refuted. Therefore, $z_i = \alpha(z_i)$ is set unconditionally.

For a fixed choice for σ and z, let $s_{\sigma,z}(\alpha)$ be the number of blocks that are good for α. For that choice, the probability that $\gamma = \alpha$ is at least

$$2^{-(kn - s_{\sigma,z}(\alpha))}$$

because for each of the kn variables $x \in X$, the probability of choosing $x = \alpha(x)$ is either

- 1, if $x = z_i$ and the other value is locally refuted (note that $b = \alpha(z_i)$ cannot be locally refuted because $(\alpha, p(\alpha))$ satisfies F)
- $1/2$ otherwise

and the first case occurs at least once for each good block. The sum of all these probabilities over $\alpha \in B$ is at most 1, so

$$\sum_{\alpha \in B} 2^{-(kn - s_{\sigma,z}(\alpha))} \leq 1$$

for every fixed choice of σ and z. Averaging over the possible choices yields

$$\begin{aligned} 1 &\geq \mathbb{E}_{\sigma,z}\left[\sum_{\alpha \in B} 2^{-(kn - s_{\sigma,z}(\alpha))}\right] \\ &= 2^{-kn} \sum_{\alpha \in B} \mathbb{E}_{\sigma,z}\left[2^{s_{\sigma,z}(\alpha)}\right]. \end{aligned}$$

For a fixed α, each block i is good with probability at least $1/k^2$, so we get

$$\mathbb{E}_{\sigma,z}\left[s_{\sigma,z}(\alpha)\right] \geq \frac{n}{k^2}.$$

Because 2^x is a convex function, we can apply Jensen's inequality and get

$$\mathbb{E}_{\sigma,z}\left[2^{s_{\sigma,z}(\alpha)}\right] \geq 2^{n/k^2}.$$

Combining these yields

$$1 \geq 2^{-kn} \cdot |B| \cdot 2^{n/k^2}$$

and therefore $|B| \leq 2^{kn - n/k^2}$. ◀

In essence, Lemma 4.6 gives an upper bound on the number of existential assignments a single response-k-CNF can cover for the k-`Equality` formulas. This can then be used to show the actual cost lower bounds.

▶ **Lemma 4.7.** *For all $k \geq 1$, $\mathrm{cost}_{k\text{-}CNF}(k\text{-}\mathtt{Equality}_n) = \exp(\Omega(n))$.*

Proof. Let $S = (S_1)$ be a k-CNF Two-Step-Strategy for k-$\mathtt{Equality}_n$, and define X, U, $p(\alpha)$ and B_F as in Lemma 4.6. For each α consider the k-CNF $F_\alpha = S_1(\alpha)$. Due to Definition 3.7 of a Two-Step-Strategy, and since p is the unique universal winning strategy of k-$\mathtt{Equality}_n$, the outputs of F_α on α are fixed:

$$\begin{aligned} &F_\alpha(\alpha, p(\alpha)) = 1, \\ &\quad F_\alpha(\alpha, \beta) = 0 \quad \text{for every } \beta \neq p(\alpha). \end{aligned}$$

Therefore, $\alpha \in B_{F_\alpha}$, so the B_F for $F \in \mathrm{rng}(S_1)$ cover all assignments $\alpha \in \langle X \rangle$.

Applying Lemma 4.6 yields

$$2^{kn} = |\langle X \rangle| \leq \sum_{F \in \mathrm{rng}(S_1)} |B_F| \leq |\mathrm{rng}(S_1)| \cdot 2^{kn - n/k^2}$$

and $|\mathrm{rng}(S_1)| \geq 2^{n/k^2}$. As there is only one universal block and $|F| = 1$ for every k-CNF F, we get $\mathrm{cost}(S) = \mathrm{size}(S_1) = |\mathrm{rng}(S_1)|$. As S was chosen arbitrarily,

$$\mathrm{cost}_{k\text{-CNF}}(k\text{-}\mathtt{Equality}_n) \geq 2^{n/k^2} = \exp(\Omega(n)).$$ ◀

We now have everything required to show the k-`Equality` lower bounds for $\mathcal{Q}$-k-DNF.

▶ **Corollary 4.8.** *For all $k \geq 1$, k-`Equality` requires exponential proofs in $\mathcal{Q}$-k-DNF.*

Proof. This follows from the size-cost-capacity theorem (Theorem 3.10) together with high $\mathrm{cost}_{k\text{-CNF}}(k\text{-}\mathtt{Equality})$ (Lemma 4.7) and the low k-CNF-capacity for all $\mathcal{Q}$-k-DNF proofs (Lemma 4.5). ◀

Using these lower bounds, we can finally separate the different levels in the $\mathcal{Q}$-k-DNF hierarchy.

▶ **Theorem 4.9.** *For every $k \in \mathbb{N}$, $\mathcal{Q}$-$(k+1)$-DNF is exponentially stronger than $\mathcal{Q}$-k-DNF.*

Proof. The simulation is syntactic, every $\mathcal{Q}$-k-DNF proof is directly a $\mathcal{Q}$-$(k+1)$-DNF proof. The separating formula is k-`Equality`. Lemma 4.3 gives a short $\mathcal{Q}$-$(k+1)$-DNF proof and Corollary 4.8 gives the lower bound for $\mathcal{Q}$-k-DNF proofs. ◀

5 A second application: PTFs

As a second application we consider degree-k threshold functions as lines.

▶ **Definition 5.1** (degree-k threshold function). *Let V be a set of Boolean variables. A function $f : \langle V \rangle \to \{0,1\}$ is a* degree-k polynomial threshold function (PTF) *if there is a polynomial p with integer coefficients, given as a sum of monomials of degree at most k, such that*

$$f(\sigma) = 1 \iff p(\sigma) \geq 0 \qquad \text{for all } \sigma \in \langle V \rangle.$$

We denote the class of such functions (together with their representations) by PTF_k.

▶ Remark 5.2 (normalisation). Since σ ranges over $\{0,1\}$-assignments, we may assume p to be multilinear ($x^2 = x$). Integer coefficients are no restriction: rational coefficients can be scaled, and real coefficients can be transformed to rationals without changing any sign on the finite domain $\langle V \rangle$. Note that for integral p we have $f(\sigma) = 0 \iff p(\sigma) \leq -1$, so PTF_k is closed under negation: $\neg(p \geq 0)$ is represented by $-p - 1 \geq 0$.

We can encode a clause C into a PTF_1 by

$$\mathrm{enc}(C) := \big(\sum_{x \in C} x + \sum_{\neg x \in C} (1 - x) - 1 \geq 0 \big).$$

The conditioning property is apparent. As such, PTF_k are base representation classes for all $k \geq 1$ and using Definition 3.2, we can define the QBF proof system $\mathcal{Q}$-PTF_k.

▶ **Lemma 5.3.** *For every $k \geq 1$, $\mathcal{Q}$-PTF_k p-simulates and is exponentially stronger than $\mathcal{Q}$-k-DNF.*

Proof. To show the p-simulation, we show that every k-DNF can be transformed into an equivalent PTF_k. Let a k-DNF $D = \bigvee_t T_t$ be given. It can be encoded as a PTF_k as follows:

$$\mathrm{enc}(D) = \big(\sum_t \prod_{x \in T_t} x \prod_{\neg x \in T_t} (1-x) - 1 \geq 0\big).$$

Since every T_t is a term with at most k literals and k is constant, writing $\mathrm{enc}(D)$ as a sum of monomials blows up the size at most polynomially (in the number of variables).

Let $\pi = (D_1, \dots, D_s)$ be a $\mathcal{Q}$-k-DNF proof of a QBF Φ. Then $\pi' = (\mathrm{enc}(D_1), \dots, \mathrm{enc}(D_s))$ is a $\mathcal{Q}$-PTF_k proof of Φ. The separating formula is Q-Maj. Q-Maj is hard for the semantic version of $\mathcal{Q}$-AC^0-Frege [7], which trivially simulates all $\mathcal{Q}$-k-DNF systems (every k-DNF is a depth-2 AC^0 formula). At the same time, it is easy for semantic QBF cutting planes $\mathcal{Q}$-$\mathrm{CP} = \mathcal{Q}$-$\mathrm{PTF}_1$ [4] and, as such, easy for all $\mathcal{Q}$-PTF_k. ◀

We use $\mathcal{L} = \mathrm{PTF}_k^{\mathrm{comp}}$. Similar to above, we will omit the superscript for better readability. As PTF_k are closed under negation, $\mathrm{capacity}_{\mathrm{PTF}_k}(\pi) = 1$ for every $\mathcal{Q}$-PTF_k proof π (we again use the convention $|f| = 1$ for every $f \in \mathrm{PTF}_k$, so that size and capacity count the number of distinct lines and responses). As such, using Theorem 3.10, every $\mathrm{cost}_{\mathrm{PTF}_k}$ lower bound for a QBF Φ directly yields a proof size lower bound for $\mathcal{Q}$-PTF_k refutations of Φ.

▶ **Corollary 5.4** (threshold cost lower bounds). *If $\mathrm{cost}_{\mathrm{PTF}_k}(\Phi) \geq s$, then every $\mathcal{Q}$-PTF_k refutation of Φ has size at least s.*

Proof. This is immediate from Theorem 3.10 and the unit capacity. ◀

We separate the first two levels of the $\mathcal{Q}$-PTF_k hierarchy:

▶ **Proposition 5.5.** *$\mathcal{Q}$-PTF_2 p-simulates and is exponentially stronger than $\mathcal{Q}$-PTF_1. A separating formula is 1-*`Equality`*.*

Proof. The simulation is syntactic. Since $\mathcal{Q}$-PTF_1 is equivalent to $\mathcal{Q}$-CP, the hardness of 1-`Equality` for $\mathcal{Q}$-PTF_1 is the result of the `Equality` lower bound of $\mathcal{Q}$-CP by Beyersdorff et al. [5]. $\mathcal{Q}$-PTF_2 has short refutations for 1-`Equality`, following its simulation of $\mathcal{Q}$-2-DNF (Lemma 5.3) and the $\mathcal{Q}$-k-DNF upper bound (Lemma 4.3). ◀

In general, we conjecture that our technique can be used to show that k-`Equality` is hard for $\mathcal{Q}$-PTF_k.

▶ **Conjecture 5.6.** $\mathrm{cost}_{PTF_k}(k\text{-}\mathtt{Equality}_n) = \exp(\Omega(n))$.

▶ **Corollary 5.7** (conditional separation of the $\mathcal{Q}$-PTF_k hierarchy). *Assuming Conjecture 5.6, $\mathcal{Q}$-PTF_{k+1} p-simulates and is exponentially stronger than $\mathcal{Q}$-PTF_k for every $k \geq 1$. In particular, the $\mathcal{Q}$-PTF_k hierarchy is exponentially separated at every level.*

Proof. Fix some arbitrary $k \geq 1$. $\mathcal{Q}$-$(k+1)$-DNF have polynomial size refutations of k-`Equality` (Lemma 4.3) and due to Lemma 5.3, so do $\mathcal{Q}$-PTF_{k+1}. Using Conjecture 5.6, Corollary 5.4 yields an exponential lower bound for $\mathcal{Q}$-PTF_k refutations for the same formulas. The simulation is syntactic. ◀

Figure 1 illustrates the simulation order of these systems, including our conjectured separations.

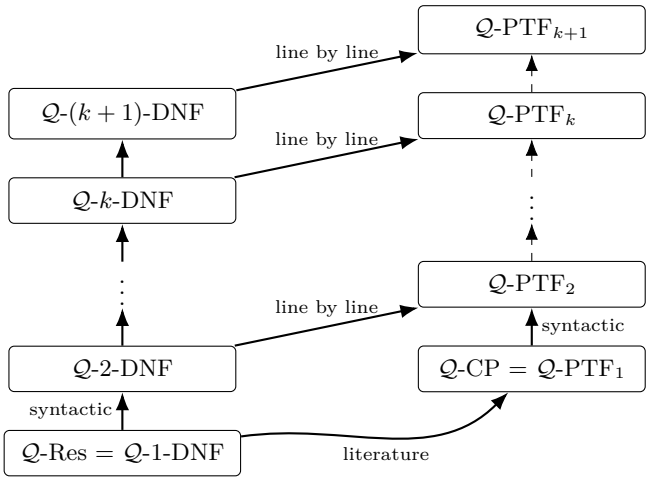

Figure 1 Simulation structure around the $\mathcal{Q}$-DNF, $\mathcal{Q}$-CP, and $\mathcal{Q}$-PTF systems. Stronger systems are placed higher whenever a simulation is shown; all arrows denote p-simulations. Dashed arrows mark simulations with conjectured separations, while solid lines show proven separations.

6 Conclusion

We introduced a fine-grained, parametrised generalisation of the size-cost-capacity technique that applies to strong semantic QBF proof systems of unbounded capacity, and used it to exponentially separate all levels of the $\mathcal{Q}$-k-DNF hierarchy.

A natural next goal is to separate the $\mathcal{Q}$-PTF$_k$ hierarchy. We conjecture that k-`Equality` has exponential PTF$_k$-cost, which by Theorem 3.10 would immediately yield exponential separations between all consecutive levels. More generally, our two-step technique invites applications to further semantic QBF systems with e.g. circuit classes as lines, such as branching programs or low-depth circuits. The challenge is to identify response classes $\mathcal{L}$ that rescale capacity to a constant while keeping the cost of natural QBF families exponential. Progress in this direction would chart the largely unexplored territory between QBF resolution and QBF Frege.

Acknowledgements

This work was partly done during a research stay at IMSc Chennai. We thank Meena Mahajan for helpful discussions on the topic of this paper.

References

1 Sanjeev Arora and Boaz Barak. *Computational Complexity – A Modern Approach*. Cambridge University Press, 2009. ISBN 978-0-521-42426-4.
2 Olaf Beyersdorff. Proof complexity of quantified Boolean logic – a survey. In Marco Benini, Olaf Beyersdorff, Michael Rathjen, and Peter Schuster, editors, *Mathematics for Computation (M4C)*, pages 353–391. World Scientific, Singapore, 2022.
3 Olaf Beyersdorff and Benjamin Böhm. Understanding the relative strength of QBF CDCL solvers and QBF resolution. *Log. Methods Comput. Sci.*, 19(2), 2023.
4 Olaf Beyersdorff, Leroy Chew, Meena Mahajan, and Anil Shukla. Understanding cutting planes for QBFs. *Inf. Comput.*, 262:141–161, 2018.

5 Olaf Beyersdorff, Joshua Blinkhorn, and Luke Hinde. Size, cost, and capacity: A semantic technique for hard random QBFs. *Logical Methods in Computer Science*, 15(1), 2019.

6 Olaf Beyersdorff, Leroy Chew, and Mikolás Janota. New resolution-based QBF calculi and their proof complexity. *ACM Transactions on Computation Theory*, 11(4):26:1–26:42, 2019.

7 Olaf Beyersdorff, Ilario Bonacina, Leroy Chew, and Jan Pich. Frege systems for quantified Boolean logic. *J. ACM*, 67(2):9:1–9:36, 2020.

8 Olaf Beyersdorff, Luke Hinde, and Ján Pich. Reasons for hardness in QBF proof systems. *ACM Transactions on Computation Theory*, 12(2):10:1–10:27, 2020.

9 Olaf Beyersdorff, Mikolás Janota, Florian Lonsing, and Martina Seidl. Quantified Boolean formulas. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, Frontiers in Artificial Intelligence and Applications, pages 1177–1221. IOS Press, 2021.

10 Olaf Beyersdorff, Joshua Blinkhorn, Meena Mahajan, and Tomás Peitl. Hardness characterisations and size-width lower bounds for QBF resolution. *ACM Trans. Comput. Log.*, 24(2):10:1–10:30, 2023.

11 Olaf Beyersdorff, Ilario Bonacina, Kaspar Kasche, Meena Mahajan, and Luc Nicolas Spachmann. Semi-algebraic proof systems for QBF. In *28th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, volume 341 of *LIPIcs*, pages 5:1–5:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi: 10.4230/LIPIcs.SAT.2025.5.

12 Sam Buss and Jakob Nordström. Proof complexity and SAT solving. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, Frontiers in Artificial Intelligence and Applications, pages 233–350. IOS Press, 2021.

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

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

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

16 Juan Luis Esteban, Nicola Galesi, and Jochen Messner. On the complexity of resolution with bounded conjunctions. *Theoretical Computer Science*, 321(2–3):347–370, 2004.

17 Enrico Giunchiglia, Paolo Marin, and Massimo Narizzano. Reasoning with quantified Boolean formulas. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 761–780. IOS Press, 2009. ISBN 978-1-58603-929-5. doi: 10.3233/978-1-58603-929-5-761. URL `https://doi.org/10.3233/978-1-58603-929-5-761`.

18 Alexandra Goultiaeva, Allen Van Gelder, and Fahiem Bacchus. A uniform approach for generating proofs and strategies for both true and false QBF formulas. In *IJCAI*, pages 546–553, 2011.

19 Mikolás Janota and Joao Marques-Silva. Expansion-based QBF solving versus Q-resolution. *Theor. Comput. Sci.*, 577:25–42, 2015.

20 Mikolás Janota, William Klieber, João Marques-Silva, and Edmund M. Clarke. Solving QBF with counterexample guided refinement. *Artif. Intell.*, 234:1–25, 2016.

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

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

23 Jan Krajíček. On the weak pigeonhole principle. *Fundamenta Mathematicae*, 170:123–140, 2001.

24 João P. Marques Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, Frontiers in Artificial Intelligence and Applications, pages 133–182. IOS Press, 2021.

25 Markus N. Rabe and Leander Tentrup. CAQE: A certifying QBF solver. In Roope Kaivola and Thomas Wahl, editors, *Formal Methods in Computer-Aided Design (FMCAD)*, pages 136–143. IEEE, 2015.

26 Agnes Schleitzer and Olaf Beyersdorff. Computationally hard problems are hard for QBF proof systems too. In *39th AAAI Conference on Artificial Intelligence*, pages 11336–11344. AAAI Press, 2025.

27 Nathan Segerlind, Samuel R. Buss, and Russell Impagliazzo. A switching lemma for small restrictions and lower bounds for k-DNF resolution. *SIAM Journal on Computing*, 33(5):1171–1200, 2004.

28 Ankit Shukla, Armin Biere, Luca Pulina, and Martina Seidl. A survey on applications of quantified Boolean formulas. In *Proc. IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 78–84, 2019.

29 Allen Van Gelder. Contributions to the theory of practical quantified Boolean formula solving. In *Proc. Principles and Practice of Constraint Programming (CP)*, pages 647–663, 2012.

30 Moshe Y. Vardi. Boolean satisfiability: theory and engineering. *Commun. ACM*, 57(3):5, 2014.

31 Lintao Zhang and Sharad Malik. Conflict driven learning in a quantified Boolean satisfiability solver. In *Proc. IEEE/ACM International Conference on Computer-aided Design (ICCAD)*, pages 442–449, 2002.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il