

Span-Program-Based Quantum Algorithm for Evaluating Unbalanced Formulas

Ben W. Reichardt^(✉)

Institute for Quantum Computing, University of Waterloo, Waterloo, Canada
breic@iqc.ca

Abstract. The formula-evaluation problem is defined recursively. A formula’s evaluation is the evaluation of a gate, the inputs of which are themselves independent formulas. Despite this pure recursive structure, the problem is combinatorially difficult for classical computers.

A quantum algorithm is given to evaluate formulas over any finite boolean gate set. Provided that the complexities of the input subformulas to any gate differ by at most a constant factor, the algorithm has optimal query complexity. After efficient preprocessing, it is nearly time optimal. The algorithm is derived using the span program framework. It corresponds to the composition of the individual span programs for each gate in the formula. Thus the algorithm’s structure reflects the formula’s recursive structure.

1 Introduction

A k -bit *gate* is a function $f : \{0, 1\}^k \rightarrow \{0, 1\}$. A *formula* φ over a set of gates \mathcal{S} is a rooted tree in which each node with k children is associated to a k -bit gate from \mathcal{S} , for $k = 1, 2, \dots$. Any such tree with n leaves naturally defines a function $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$, by placing the input bits on the leaves in a fixed order and evaluating the gates recursively toward the root. Such functions are often called *read-once* formulas, as each input bit is associated to one leaf only.

The formula-evaluation problem is to evaluate a formula φ over \mathcal{S} on an input $x \in \{0, 1\}^n$. The formula is given, but the input string x must be queried one bit at a time. How many queries to x are needed to compute $\varphi(x)$? We would like to understand this complexity as a function of \mathcal{S} and asymptotic properties of φ . Roughly, larger gate sets allow φ to have less structure, which increases the complexity of evaluating φ . Another important factor is often the balancedness of the tree φ . Unbalanced formulas often seem to be more difficult to evaluate.

For applications, the most important gate set consists of all AND and OR gates. Formulas over this set are known as AND-OR formulas. Evaluating such a formula solves the decision version of a MIN-MAX tree, also known as a two-player game tree. Unfortunately, the complexity of evaluating formulas, even over this limited gate set, is unknown, although important special cases have been solved. The problem over much larger gate sets appears to be combinatorially intractable. For some formulas, it is known that “non-directional” algorithms

that do not work recursively on the structure of the formula perform better than any recursive procedure.

In this article, we show that the formula-evaluation problem becomes dramatically simpler when we allow the algorithm to be a bounded-error quantum algorithm, and allow it coherent query access to the input string x . Fix \mathcal{S} to be any finite set of gates. We give an optimal quantum algorithm for evaluating “almost-balanced” formulas over \mathcal{S} . The balance condition states that the complexities of the input subformulas to any gate differ by at most a constant factor, where complexity is measured by the general adversary bound Adv^\pm . In general, Adv^\pm is the value of an exponentially large semi-definite program (SDP). For a formula φ with constant-size gates, though, $\text{Adv}^\pm(\varphi)$ can be computed efficiently by solving constant-size SDPs for each gate.

To place this work in context, some classical and quantum results for evaluating formulas are summarized in Table 1. The stated upper bounds are on query complexity and not time complexity. However, for the OR_n and balanced $\text{AND}_2\text{-OR}_2$ formulas, the quantum algorithms’ running times are only slower by a poly-logarithmic factor. For the other formulas, the quantum algorithms’ running times are slower by a poly-logarithmic factor provided that:

1. A polynomial-time classical preprocessing step, outputting a string $s(\varphi)$, is not charged for.
2. The algorithms are allowed unit-cost coherent access to $s(\varphi)$.

Our algorithm is based on the framework relating span programs and quantum algorithms from [Rei09]. Previous work has used span programs to develop quantum algorithms for evaluating formulas [RŠ08]. Using this and the observation that the optimal span program witness size for a boolean function f equals the general adversary bound $\text{Adv}^\pm(f)$, Ref. [Rei09] gives an optimal quantum algorithm for evaluating “adversary-balanced” formulas over an arbitrary finite gate set. The balance condition is that each gate’s input subformulas have equal general adversary bounds.

In order to relax this strict balance requirement, we must maintain better control in the recursive analysis. To help do so, we define a new span program complexity measure, the “full witness size.” This complexity measure has implications for developing time- and query-efficient quantum algorithms based on span programs. Essentially, using a second result from [Rei09], that properties of eigenvalue-zero eigenvectors of certain bipartite graphs imply “effective” spectral gaps around zero, it allows quantum algorithms to be based on span programs with free inputs. This can simplify the implementation of a quantum walk on the corresponding graph.

Besides allowing a relaxed balance requirement, our approach has the additional advantage of making the constants hidden in the big- O notation more explicit. The formula-evaluation quantum algorithms in [RŠ08, Rei09] evaluate certain formulas φ using $O(\text{Adv}^\pm(\varphi))$ queries, where the hidden constant depends on the gates in \mathcal{S} in a complicated manner. It is not known how to upper-bound the hidden constant in terms of, say, the maximum fan-in k of a

Table 1. Comparison of some classical and quantum query complexity results for formula evaluation. Here \mathcal{S} is any fixed, finite gate set, and the exponent α is given by $\alpha = \log_2\left(\frac{1+\sqrt{33}}{4}\right) \approx 0.753$. Under certain assumptions, the algorithms' running times are only poly-logarithmically slower.

Formula φ	Randomized, zero-error query complexity $R(\varphi)$	Quantum bounded-error query complexity $Q(\varphi)$
OR_n	n	$\Theta(\sqrt{n})$ [Gro96, BBBV97]
Balanced AND_2 - OR_2	$\Theta(n^\alpha)$ [SW86]	$\Theta(\sqrt{n})$ [FGG08, ACR ⁺ 10]
Well-balanced AND - OR	Tight recursion [SW86]	$\Theta(\sqrt{n})$ [ACR ⁺ 10], (Theorem 8)
Approx.-balanced AND - OR		$\Omega(\sqrt{n})$ [BS04]
Arbitrary AND - OR	$\Omega(n^{0.51})$ [HW91]	$O(\sqrt{n} \log n)$ [Rei11]
Balanced MAJ_3 ($n = 3^d$)	$\Omega((7/3)^d), O(2.654^d)$ [JKS03]	$\Theta(2^d)$ [RŠ08]
Balanced over \mathcal{S}		$\Theta(\text{Adv}^\pm(\varphi))$ [Rei09]
Almost-balanced over \mathcal{S}		$\Theta(\text{Adv}^\pm(\varphi))$ (Theorem 7)

gate in \mathcal{S} . In contrast, the approach we follow here allows bounding this constant by an exponential in k .

It is known that the general adversary bound is a nearly tight lower bound on quantum query complexity for *any* boolean function [Rei09], including in particular boolean formulas. However, this comes with no guarantees on time complexity. The main contribution of this paper is to give a nearly time-optimal algorithm for formula evaluation. The algorithm is also tight for query complexity, removing the extra logarithmic factor from the bound in [Rei09].

Additionally, we apply the same technique to study AND - OR formulas. For this special case, special properties of span programs for AND and for OR gates allow the almost-balance condition to be significantly weakened. Ambainis et al. [ACR⁺10] have studied this case previously. By applying the span program framework, we identify a slight weakness in their analysis. Tightening the analysis extends the algorithm's applicability to a broader class of AND - OR formulas.

A companion paper [Rei11] applies the span program framework to the problem of evaluating *arbitrary* AND - OR formulas. By studying the full witness size for span programs constructed using a novel composition method, it gives an $O(\sqrt{n} \log n)$ -query quantum algorithm to evaluate a formula of size n , for which the time complexity is poly-logarithmically worse after preprocessing. This nearly matches the $\Omega(\sqrt{n})$ lower bound, and improves a $\sqrt{n}2^{O(\sqrt{\log n})}$ -query quantum algorithm from [ACR⁺10]. Reference [Rei11] shares the broader motivation of this paper, to study span program properties and design techniques that lead to time-efficient quantum algorithms.

Sections 1.1 and 1.2 below give further background on the formula-evaluation problem, for classical and quantum algorithms. Section 1.3 precisely states our main theorem, the proof of which is given in Sect. 3 after some background on

span programs. The theorem for approximately balanced AND-OR formulas is stated in Sect. 1.4, and proved in Sect. 4.

1.1 History of the Formula-Evaluation Problem for Classical Algorithms

For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let $D(f)$ be the least number of input bit queries sufficient to evaluate f on any input with zero error. $D(f)$ is known as the deterministic decision-tree complexity of f , or the deterministic query complexity of f . Let the randomized decision-tree complexity of f , $R(f) \leq D(f)$, be the least *expected* number of queries required to evaluate f with zero error (i.e., by a Las Vegas randomized algorithm). Let the Monte Carlo decision-tree complexity, $R_2(f) = O(R(f))$, be the least number of queries required to evaluate f with error probability at most $1/3$ (i.e., by a Monte Carlo randomized algorithm).

Classically, formulas over the gate set $\mathcal{S} = \{\text{NAND}_k : k \in \mathbf{N}\}$ have been studied most extensively, where $\text{NAND}_k(x_1, \dots, x_k) = 1 - \prod_{j=1}^k x_j$. By De Morgan's rules, any formula over NAND gates can also be written as a formula in which the gates at an even distance from the formula's root are AND gates and those at an odd distance away are OR gates, with some inputs or the output possibly complemented. Thus formulas over \mathcal{S} are also known as AND-OR formulas.

For any AND-OR formula φ of size n , i.e., on n inputs, $D(\varphi) = n$. However, randomization gives a strict advantage; $R(\varphi)$ and $R_2(\varphi)$ can be strictly smaller. Indeed, let φ_d be the complete, binary AND-OR formula of depth d , corresponding to the tree in which each internal vertex has two children and every leaf is at distance d from the root, with alternating levels of AND and OR gates. Its size is $n = 2^d$. Snir [Sni85] has given a randomized algorithm for evaluating φ_d using in expectation $O(n^\alpha)$ queries, where $\alpha = \log_2\left(\frac{1+\sqrt{33}}{4}\right) \approx 0.753$ [SW86]. This algorithm, known as randomized alpha-beta pruning, evaluates a random subformula recursively, and only evaluates the second subformula if necessary. Saks and Wigderson [SW86] have given a matching lower bound on $R(\varphi_d)$, which Santha has extended to hold for Monte Carlo algorithms, $R_2(\varphi_d) = \Omega(n^\alpha)$ [San95].

Thus the query complexities have been characterized for the complete, binary AND-OR formulas. In fact, the tight characterization works for a larger class of formulas, called "well balanced" formulas by [San95]. This class includes, for example, alternating $\text{AND}_2\text{-OR}_2$ formulas where for some d every leaf is at depth d or $d - 1$, Fibonacci trees and binomial trees [SW86]. It also includes skew trees, for which the depth is the maximal $n - 1$.

For arbitrary AND-OR formulas, on the other hand, little is known. It has been conjectured that complete, binary AND-OR formulas are the easiest to evaluate, and that in particular $R(\varphi) = \Omega(n^\alpha)$ for any size- n AND-OR formula φ [SW86]. However, the best general lower bound is $R(\varphi) = \Omega(n^{0.51})$, due to Heiman and Wigderson [HW91]. Reference [HW91] also extends the result of [SW86] to allow for AND and OR gates with fan-in more than two.

It is perhaps not surprising that formulas over most other gate sets \mathcal{S} are even less well understood. For example, Boppana has asked the complexity of

evaluating the complete ternary majority (MAJ₃) formula of depth d [SW86], and the best published bounds on its query complexity are $\Omega((7/3)^d)$ and $O((2.6537\dots)^d)$ [JKS03]. In particular, the naïve, “directional,” generalization of the randomized alpha-beta pruning algorithm is to evaluate recursively two random immediate subformulas and, if they disagree, then also the third. This algorithm uses $O((8/3)^d)$ expected queries, and is suboptimal. This suggests that the complete MAJ₃ formulas are significantly different from the complete AND-OR formulas.

Heiman, Newman and Wigderson have considered read-once threshold formulas in an attempt to separate the complexity classes TC⁰ from NC¹ [HNW93]. That is, they allow the gate set to be the set of Hamming-weight threshold gates $\{T_m^k : m, k \in \mathbf{N}\}$ defined by $T_m^k : \{0, 1\}^k \rightarrow \{0, 1\}$, $T_m^k(x) = 1$ if and only if the Hamming weight of x is at least m . AND, OR and majority gates are all special cases of threshold gates. Heiman et al. prove that $R(\varphi) \geq n/2^d$ for φ a threshold formula of depth d , and in fact their proof extends to gate sets in which every gate “contains a flip” [HNW93]. This implies that a large depth is necessary for the randomized complexity to be much lower than the deterministic complexity.

Of course there are some trivial gate sets for which the query complexity is fully understood, for example, the set of parity gates. Overall, though, there are many more open problems than results. Despite its structure, formula evaluation appears to be combinatorially complicated. However, there is another approach, to try to leverage the power of quantum computers. Surprisingly, the formula-evaluation problem simplifies considerably in this different model of computation.

1.2 History of the Formula-Evaluation Problem for Quantum Algorithms

In the quantum query model, the input bits can be queried coherently. That is, the quantum algorithm is allowed unit-cost access to the unitary operator O_x , called the input oracle, defined by

$$O_x : |\varphi\rangle \otimes |j\rangle \otimes |b\rangle \mapsto |\varphi\rangle \otimes |j\rangle \otimes |b \oplus x_j\rangle . \quad (1.1)$$

Here $|\varphi\rangle$ is an arbitrary pure state, $\{|j\rangle : j = 1, 2, \dots, n\}$ is an orthonormal basis for \mathbf{C}^n , $\{|b\rangle : b = 0, 1\}$ is an orthonormal basis for \mathbf{C}^2 , and \oplus denotes addition mod two. O_x can be implemented efficiently on a quantum computer given a classical circuit that computes the function $j \mapsto x_j$ [NC00]. For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let $Q(f)$ be the number of input queries required to evaluate f with error probability at most $1/3$. It is immediate that $Q(f) \leq R_2(f)$.

Research on the formula-evaluation problem in the quantum model began with the n -bit OR function, OR _{n} . Grover gave a quantum algorithm for evaluating OR _{n} with bounded one-sided error using $O(\sqrt{n})$ oracle queries and $O(\sqrt{n} \log \log n)$ time [Gro96, Gro02]. In the classical case, on the other hand, it is obvious that $R_2(\text{OR}_n)$, $R(\text{OR}_n)$ and $D(\text{OR}_n)$ are all $\Theta(n)$.

Grover’s algorithm can be applied recursively to speed up the evaluation of more general AND-OR formulas. Call a formula *layered* if the gates at the

same depth are the same. Buhrman, Cleve and Wigderson show that a layered, depth- d , size- n AND-OR formula can be evaluated using $O(\sqrt{n} \log^{d-1} n)$ queries [BCW98]. The logarithmic factors come from using repetition at each level to reduce the error probability from a constant to be polynomially small.

Høyer, Mosca and de Wolf [HMW03] consider the case of a unitary input oracle \tilde{O}_x that maps

$$\tilde{O}_x : |\varphi\rangle \otimes |j\rangle \otimes |b\rangle \otimes |0\rangle \mapsto |\varphi\rangle \otimes |j\rangle \otimes (|b \oplus x_j\rangle \otimes |\psi_{x,j,x_j}\rangle + |b \oplus \bar{x}_j\rangle \otimes |\psi_{x,j,\bar{x}_j}\rangle), \quad (1.2)$$

where $|\psi_{x,j,x_j}\rangle, |\psi_{x,j,\bar{x}_j}\rangle$ are pure states with $\| |\psi_{x,j,x_j}\rangle \|^2 \geq 2/3$. Such an oracle can be implemented when the function $j \mapsto x_j$ is computed by a bounded-error, randomized subroutine. Høyer et al. allow access to \tilde{O}_x and \tilde{O}_x^{-1} , both at unit cost, and show that OR_n can still be evaluated using $O(\sqrt{n})$ queries. This *robustness* result implies that the $\log n$ steps of repetition used by [BCW98] are not necessary, and a depth- d layered AND-OR formula can be computed in $O(\sqrt{n} c^{d-1})$ queries, for some constant $c > 1000$. If the depth is constant, this gives an $O(\sqrt{n})$ -query quantum algorithm, but the result is not useful for the complete, binary AND-OR formula, for which $d = \log_2 n$.

In 2007, Farhi, Goldstone and Gutmann presented a quantum algorithm for evaluating complete, binary AND-OR formulas [FGG08]. Their breakthrough algorithm is not based on iterating Grover’s algorithm in any way, but instead runs a quantum walk—analogueous to a classical random walk—on a graph based on the formula. The algorithm runs in time $O(\sqrt{n})$ in a certain continuous-time query model.

Ambainis et al. discretized the [FGG08] algorithm by reinterpreting a correspondence between (discrete-time) random and quantum walks due to Szegedy [Sze04] as a correspondence between continuous-time and discrete-time quantum walks [ACR+10]. Applying this correspondence to quantum walks on certain *weighted* graphs, they gave an $O(\sqrt{n})$ -query quantum algorithm for evaluating “approximately balanced” AND-OR formulas. For example, $\text{MAJ}_3(x_1, x_2, x_3) = (x_1 \wedge x_2) \vee ((x_1 \vee x_2) \wedge x_3)$, so there is a size- 5^d AND-OR formula that computes MAJ_3^d (the complete ternary majority formula of depth d). Since the formula is approximately balanced, $Q(\text{MAJ}_3^d) = O(\sqrt{5^d})$, better than the $\Omega((7/3)^d)$ classical lower bound.

The [ACR+10] algorithm also applies to arbitrary AND-OR formulas. If φ has size n and depth d , then the algorithm, applied directly, evaluates φ using $O(\sqrt{n} d)$ queries.¹ This can be as bad as $O(n^{3/2})$ if the depth is $d = n$. However, Bshouty, Cleve and Eberly have given a formula rebalancing procedure that takes AND-OR formula φ as input and outputs an equivalent AND-OR formula φ' with depth $d' = 2^{O(\sqrt{\log n})}$ and size $n' = n 2^{O(\sqrt{\log n})}$ [BCE91, BB94]. The formula φ' can then be evaluated using $O(\sqrt{n'} d') = \sqrt{n} 2^{O(\sqrt{\log n})}$ queries.

¹ Actually, [ACR+10, Section 7] only shows a bound of $O(\sqrt{n} d^{3/2})$ queries, but this can be improved to $O(\sqrt{n} d)$ using the bounds on $\sigma_{\pm}(\varphi)$ below [ACR+10, Definition 1].

Our understanding of *lower* bounds for the formula-evaluation problem progressed in parallel to this progress on quantum algorithms. There are essentially two techniques, the *polynomial* and *adversary* methods, for lower-bounding quantum query complexity.

- The polynomial method, introduced in the quantum setting by Beals et al. [BBC⁺01], is based on the observation that after making q oracle O_x queries, the probability of any measurement result is a polynomial of degree at most $2q$ in the variables x_j .
- Ambainis generalized the classical hybrid argument, to consider the system’s entanglement when run on a superposition of inputs [Amb02]. A number of variants of Ambainis’s bound were soon discovered, including weighted versions [HNS02, BS04, Amb06, Zha05], a spectral version [BSS03], and a version based on Kolmogorov complexity [LM04]. These variants can be asymptotically stronger than Ambainis’s original unweighted bound, but are equivalent to each other [ŠS06]. We therefore term it simply “the adversary bound,” denoted by Adv.

The adversary bound is well-suited for lower-bounding the quantum query complexity for evaluating formulas. For example, Barnum and Saks proved that for any size- n AND-OR formula φ , $\text{Adv}(\varphi) = \sqrt{n}$, implying the lower bound $Q(\varphi) = \Omega(\sqrt{n})$ [BS04]. Thus the [ACR⁺10] algorithm is optimal for approximately balanced AND-OR formulas, and is nearly optimal for arbitrary AND-OR formulas. This is a considerably more complete solution than is known classically.

It is then natural to consider formulas over larger gate sets. The adversary bound continues to work well, because it transforms nicely under function composition:

Theorem 1 (Adversary bound composition [Amb06, LLS06, HLŠ05]). *Let $f : \{0, 1\}^k \rightarrow \{0, 1\}$ and let $f_j : \{0, 1\}^{m_j} \rightarrow \{0, 1\}$ for $j = 1, 2, \dots, k$. Define $g : \{0, 1\}^{m_1} \times \dots \times \{0, 1\}^{m_k} \rightarrow \{0, 1\}$ by $g(x) = f(f_1(x_1), \dots, f_k(x_k))$. Let $s = (\text{Adv}(f_1), \dots, \text{Adv}(f_k))$. Then*

$$\text{Adv}(g) = \text{Adv}_s(f). \tag{1.3}$$

See Definition 3 for the definition of the adversary bound with “costs,” Adv_s . The Adv bound equals Adv_s with uniform, unit costs $s = \vec{1}$. For a function f , $\text{Adv}(f)$ can be computed using a semi-definite program in time polynomial in the size of f ’s truth table. Therefore, Theorem 1 gives a polynomial-time procedure for computing the adversary bound for a formula φ over an arbitrary finite gate set: compute the bounds for subformulas, moving from the leaves toward the root. At an internal node f , having computed the adversary bounds for the input subformulas f_1, \dots, f_k , Eq. (1.3) says that the adversary bound for g , the subformula rooted at f , equals the adversary bound for the *gate* f with certain costs. Computing this requires $2^{O(k)}$ time, which is a constant if $k = O(1)$. For example, if f is an OR_k or AND_k gate, then $\text{Adv}_{(s_1, \dots, s_k)}(f) = \sqrt{\sum_j s_j^2}$, from

which follows immediately the [BS04] result $\text{Adv}(\varphi) = \sqrt{n}$ for a size- n AND-OR formula φ .

A special case of Theorem 1 is when the functions f_j all have equal adversary bounds, so $\text{Adv}(g) = \text{Adv}(f)\text{Adv}(f_1)$. In particular, for a function $f : \{0, 1\}^k \rightarrow \{0, 1\}$ and a natural number $d \in \mathbf{N}$, let $f^d : \{0, 1\}^{k^d} \rightarrow \{0, 1\}$ denote the complete, depth- d formula over f . That is, $f^1 = f$ and $f^d(x) = f(f^{d-1}(x_1, \dots, x_{k^{d-1}}), \dots, f^{d-1}(x_{k^{d-k^{d-1}+1}}, \dots, x_{k^d}))$ for $d > 1$. Then we obtain:

Corollary 1. For any function $f : \{0, 1\}^k \rightarrow \{0, 1\}$,

$$\text{Adv}(f^d) = \text{Adv}(f)^d . \quad (1.4)$$

In particular, Ambainis defined a boolean function $f : \{0, 1\}^4 \rightarrow \{0, 1\}$ that can be represented exactly by a polynomial of degree two, but for which $\text{Adv}(f) = 5/2$ [Amb06]. Thus f^d can be represented exactly by a polynomial of degree 2^d , but by Corollary 1, $\text{Adv}(f^d) = (5/2)^d$. For this function, the adversary bound is strictly stronger than any bound obtainable using the polynomial method. Many similar examples are given in [HLŠ06]. However, for other functions, the adversary bound is asymptotically worse than the polynomial method [ŠS06, AS04, Amb05].

In 2007, though, Høyer et al. discovered a strict generalization of Adv that also lower-bounds quantum query complexity [HLŠ06]. We call this new bound the *general adversary bound*, or Adv^\pm . For example, for Ambainis's four-bit function f , $\text{Adv}^\pm(f) \geq 2.51$ [HLŠ06]. Like the adversary bound, $\text{ADV}_s^\pm(f)$ can be computed in time polynomial in the size of f 's truth table, and also composes nicely:

Theorem 2 ([HLŠ07, Rei09]). *Under the conditions of Theorem 1,*

$$\text{Adv}^\pm(g) = \text{ADV}_s^\pm(f) . \quad (1.5)$$

In particular, if $\text{Adv}^\pm(f_1) = \dots = \text{Adv}^\pm(f_k)$, then we have $\text{Adv}^\pm(g) = \text{Adv}^\pm(f) \text{Adv}^\pm(f_1)$.

Define a formula φ to be *adversary balanced* if at each internal node, the general adversary bounds of the input subformulas are equal. In particular, by Theorem 2 this implies that $\text{Adv}^\pm(\varphi)$ is equal to the product of the general adversary bounds of the gates along any path from the root to a leaf. Complete, layered formulas are an example of adversary-balanced formulas.

Returning to upper bounds, Reichardt and Špalek [RŠ08] generalized the algorithmic approach started by [FGG08]. They gave an optimal quantum algorithm for evaluating adversary-balanced formulas over a considerably extended gate set, including in particular all functions $\{0, 1\}^k \rightarrow \{0, 1\}$ for $k \leq 3$, 69 inequivalent four-bit functions, and the gates AND_k , OR_k , PARITY_k and EQUAL_k , for $k = O(1)$. For example, $Q(\text{MAJ}_3^d) = \Theta(2^d)$.

The [RŠ08] result follows from a framework for developing formula-evaluation quantum algorithms based on *span programs*. A span program, introduced by Karchmer and Wigderson [KW93], is a certain linear-algebraic way of defining a function, which corresponds closely to eigenvalue-zero eigenvectors of certain bipartite graphs. [RŠ08] derived a quantum algorithm for evaluating certain concatenated span programs, with a query complexity upper-bounded by the span program *witness size*, denoted wsize . In particular, a special case of [RŠ08, Theorem 4.7] is:

Theorem 3 ([RŠ08]). *Fix a function $f : \{0, 1\}^k \rightarrow \{0, 1\}$. If span program P computes f , then*

$$Q(f^d) = O(\text{wsize}(P)^d) . \quad (1.6)$$

From Theorem 2, this result is optimal if $\text{wsize}(P) = \text{Adv}^\pm(f)$. The question therefore becomes how to find optimal span programs. Using an ad hoc search, [RŠ08] found optimal span programs for a variety of functions with $\text{Adv}^\pm = \text{Adv}$. Further work automated the search, by giving a semi-definite program (SDP) for the optimal span program witness size for any given function [Rei09]. Remarkably, the SDP's value always equals the general adversary bound:

Theorem 4 ([Rei09]). *For any function $f : \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$\inf_P \text{wsize}(P) = \text{Adv}^\pm(f) , \quad (1.7)$$

where the infimum is over span programs P computing f . Moreover, this infimum is achieved.

This result greatly extends the gate set over which the formula-evaluation algorithm of [RŠ08] works optimally. For example, combined with Theorem 3, it implies that $\lim_{d \rightarrow \infty} Q(f^d)^{1/d} = \text{Adv}^\pm(f)$ for every boolean function f . More generally, Theorem 4 allows the [RŠ08] algorithm to be run on formulas over any finite gate set \mathcal{S} . A factor is lost that depends on the gates in \mathcal{S} , but it will be a constant for \mathcal{S} finite. Combining Theorem 4 with [RŠ08, Theorem 4.7] gives:

Theorem 5 ([Rei09]). *Let \mathcal{S} be a finite set of gates. Then there exists a quantum algorithm that evaluates an adversary-balanced formula φ over \mathcal{S} using $O(\text{Adv}^\pm(\varphi))$ input queries. After efficient classical preprocessing independent of the input x , and assuming unit-time coherent access to the preprocessed classical string, the running time of the algorithm is $\text{Adv}^\pm(\varphi)(\log \text{Adv}^\pm(\varphi))^{O(1)}$.*

In the discussion so far, we have for simplicity focused on query complexity. The query complexity is an information-theoretic quantity that does not charge for operations independent of the input string, even though these operations may require many elementary gates to implement. For practical algorithms, it is important to be able to bound the algorithm's *running time*, which counts the cost of implementing the input-independent operations. Theorem 5 puts an optimal bound on the query complexity, and also puts a nearly optimal bound

on the algorithm's time complexity. In fact, all of the query-optimal algorithms so far discussed are also nearly time optimal.

In general, though, an upper bound on the query complexity does not imply an upper bound on the time complexity. Reference [Rei09] also generalized the span program framework of [RS08] to apply to quantum algorithms not based on formulas. The main result of [Rei09] is:

Theorem 6 ([Rei09]). *For any function $f : \mathcal{D} \rightarrow \{1, 2, \dots, m\}$, with $\mathcal{D} \subseteq \{0, 1\}^n$, $Q(f)$ satisfies*

$$Q(f) = \Omega(\text{Adv}^\pm(f)) \tag{1.8}$$

$$\text{and } Q(f) = O\left(\text{Adv}^\pm(f) \frac{\log \text{Adv}^\pm(f)}{\log \log \text{Adv}^\pm(f)} \log(m) \log \log m\right). \tag{1.9}$$

Theorem 6 in particular allows us to compute the query complexity of formulas, up to the logarithmic factor. It does *not* give any guarantees on running time. However, the analysis required to prove Theorem 6 also leads to significantly simpler proofs of Theorem 5 and the AND-OR formula results of [ACR⁺10, FGG08]. Moreover, we will see that it allows the formula-evaluation algorithms to be extended to formulas that are not adversary balanced.

1.3 Quantum Algorithm for Evaluating Almost-Balanced Formulas

We give a formula-evaluation algorithm that is both query-optimal, without a logarithmic overhead, and, after an efficient preprocessing step, nearly time optimal. Define almost balance as follows:

Definition 1. *Consider a formula φ over a gate set \mathcal{S} . For a vertex v in the corresponding tree, let φ_v denote the subformula of φ rooted at v , and, if v is an internal vertex, let g_v be the corresponding gate. The formula φ is β -balanced if for every vertex v , with children c_1, c_2, \dots, c_k ,*

$$\frac{\max_j \text{Adv}^\pm(\varphi_{c_j})}{\min_j \text{Adv}^\pm(\varphi_{c_j})} \leq \beta . \tag{1.10}$$

(If c_j is a leaf, $\text{Adv}^\pm(\varphi_{c_j}) = 1$.) *Formula φ is almost balanced if it is β -balanced for some $\beta = O(1)$.*

In particular, an adversary-balanced formula is 1-balanced. We will show:

Theorem 7. *Let \mathcal{S} be a fixed, finite set of gates. Then there exists a quantum algorithm that evaluates an almost-balanced formula φ over \mathcal{S} using $O(\text{Adv}^\pm(\varphi))$ input queries. After polynomial-time classical preprocessing independent of the input, and assuming unit-time coherent access to the preprocessed string, the running time of the algorithm is $\text{Adv}^\pm(\varphi)(\log \text{Adv}^\pm(\varphi))^{O(1)}$.*

Theorem 7 is significantly stronger than Theorem 5, which requires exact balance. There are important classes of exactly balanced formulas, such as complete, layered formulas. In fact, it is sufficient that the multiset of gates along the simple path from the root to a leaf not depend on the leaf. Moreover, sometimes different gates have the same Adv^\pm bound; see [HLŠ06] for examples. Even still, exact adversary balance is a very strict condition.

The proof of Theorem 7 is based on the span program framework developed in Ref. [Rei09]. In particular, [Rei09, Theorem 9.1] gives *two* quantum algorithms for evaluating span programs. The first algorithm is based on a discrete-time simulation of a continuous-time quantum walk. It applies to arbitrary span programs, and is used, in combination with Theorem 4, to prove Theorem 6. However, the simulation incurs a logarithmic query overhead and potentially worse time complexity overhead, so this algorithm is not suitable for proving Theorem 7.

The second algorithm in [Rei09] is based directly on a discrete-time quantum walk, similar to previous optimal formula-evaluation algorithms [ACR⁺10, RŠ08]. However, this algorithm does not apply to an arbitrary span program. A bound is needed on the operator norm of the entry-wise absolute value of the weighted adjacency matrix for a corresponding graph. Further graph sparsity conditions are needed for the algorithm to be time efficient (see Theorem 9).

Unfortunately, the span program from Theorem 4 will not generally satisfy these conditions. Theorem 4 gives a *canonical* span program ([Rei09, Definition 5.1]). Even for a simple formula, the optimal canonical span program will typically correspond to a dense graph with large norm.

An example should clarify the problem. Consider the AND-OR formula $\psi(x) = ((x_1 \wedge x_2) \vee x_3) \wedge x_4 \vee (x_5 \wedge [x_6 \vee x_7])$, and consider the two graphs in Fig. 1. For an input $x \in \{0, 1\}^7$, modify the graphs by attaching dangling edges to every vertex j for which $x_j = 0$. Observe then that each graph has an eigenvalue-zero eigenvector supported on vertex 0—called a *witness*—if and only if $\psi(x) = 1$. The graphs correspond to different span programs computing ψ , and the quantum algorithm works essentially by running a quantum walk starting at vertex 0 in order to detect the witness. The graph on the left is a significantly simplified version of a canonical span program for ψ , and its density still makes it difficult to implement the quantum walk.

We will be guided by the second, simpler graph. Instead of applying Theorem 4 to φ as a whole, we apply it separately to every gate in the formula. We then compose these span programs, one per gate, according to the formula, using *direct-sum composition* (Definition 6). In terms of graphs, direct-sum composition attaches the output vertex of one span program’s graph to an input vertex of the next [RŠ08]. This leads to a graph whose structure somewhat follows the structure of the formula φ , as the graph in Fig. 1(b) follows the structure of ψ . (However, the general case will be more complicated than shown, as we are plugging together constant-size graph gadgets, and there may be duplication of some subgraphs.)




Fig. 1. Graphs corresponding to two span programs both computing the same function.

Direct-sum composition keeps the maximum degree and norm of the graph under control—each is at most twice its value for the worst single gate. Therefore the second [Rei09] algorithm applies. However, direct-sum composition also leads to additional overhead. In particular, a witness in the first graph will be supported only on numbered vertices (note that the graph is bipartite), whereas a witness in the second graph will be supported on some of the internal vertices as well. This means roughly that the second witness will be harder to detect, because after normalization its overlap on vertex 0 will be smaller. Scale both witnesses so that the amplitude on vertex 0 is one. The *witness size* (wsize) measures the squared length of the witness only on numbered vertices, whereas the *full witness size* (fwsiz) measures the squared length on all vertices. For [Rei09], it was sufficient to consider only span program witness size, because for canonical span programs like in Fig. 1(a) the two measures are equal. (For technical reasons, we will actually define fwsiz to be $1 + \text{wsize}$ even in this case.) For our analysis, we will need to bound the full witness size in terms of the witness size. We maintain this bound in a recursion from the formula’s leaves toward its root.

A span program is called *strict* if every vertex on one half of the bipartite graph is either an input vertex (vertices 1–7 in the graphs of Fig. 1) or the output vertex (vertex 0). Thus the first graph in the example above corresponds to a strict span program, and the second does not. The original definition of span programs, in [KW93], allowed for only strict span programs. This was sensible because any other vertices on the input/output part of the graph’s bipartition can always be projected away, yielding a strict span program that computes the same function. For developing time-efficient quantum algorithms, though, it seems important to consider span programs that are not strict. Unfortunately, going backwards, e.g., from 1(a) to 1(b), is probably difficult in general.

Theorem 7 does *not* follow from the formula-evaluation techniques of [RŠ08], together with Theorem 3 from [Rei09]. This tempting approach falls into intractable technical difficulties. In particular, the same span program can be used at two vertices v and w in φ only if $g_v = g_w$ and the general adversary bounds of v 's input subformulas are the same as those for w 's inputs up to simultaneous scaling. In general, then, an almost-balanced formula will require an unbounded number of different span programs. However, the analysis in [RŠ08] loses a factor that depends badly on the individual span programs. Since the dependence is not continuous, even showing that the span programs in use all lie within a compact set would not be sufficient to obtain an $O(1)$ upper bound. In contrast, the approach we follow here allows bounding the lost factor by an exponential in k , uniformly over different gate imbalances.

1.4 Quantum Algorithm to Evaluate Approximately Balanced AND-OR Formulas

Ambainis et al. [ACR⁺10] use a weaker balance criterion for AND-OR formulas than Definition 1. They define an AND-OR formula to be approximately balanced if $\sigma_-(\varphi) = O(1)$ and $\sigma_+(\varphi) = O(n)$. Here n is the size of the formula, i.e., the number of leaves, and $\sigma_-(\varphi)$ and $\sigma_+(\varphi)$ are defined by:

Definition 2. For each vertex v in a formula φ , let

$$\begin{aligned}\sigma_-(v) &= \max_{\xi} \sum_{w \in \xi} \frac{1}{\text{Adv}^{\pm}(\varphi_w)} \\ \sigma_+(v) &= \max_{\xi} \sum_{w \in \xi} \text{Adv}^{\pm}(\varphi_w)^2 \quad ,\end{aligned}\tag{1.11}$$

with each maximum taken over all simple paths ξ from v to a leaf. Let $\sigma_{\pm}(\varphi) = \sigma_{\pm}(r)$, where r is the root of φ .

Recall that $\text{Adv}^{\pm}(\varphi) = \text{Adv}(\varphi) = \sqrt{n}$ for an AND-OR formula. Definition 1 is a stricter balance criterion because β -balance of a formula φ implies (by Lemma 3) that $\sigma_-(\varphi)$ and $\sigma_+(\varphi)$ are both dominated by geometric series. However, the same steps followed by the proof of Theorem 7 still suffice for proving the [ACR⁺10] result, and, in fact, for strengthening it. We show:

Theorem 8. Let φ be an AND-OR formula of size n . Then after polynomial-time classical preprocessing that does not depend on the input x , $\varphi(x)$ can be evaluated by a quantum algorithm with error at most $1/3$ using $O(\sqrt{n} \sigma_-(\varphi))$ input queries. The algorithm's running time is $\sqrt{n} \sigma_-(\varphi) (\log n)^{O(1)}$ assuming unit-cost coherent access to the preprocessed string.

For the special case of AND-OR formulas with $\sigma_-(\varphi) = O(1)$, Theorem 8 strengthens Theorem 7. The requirement that $\sigma_-(\varphi) = O(1)$ allows for some gates in the formula to be very unbalanced. Theorem 8 also strengthens

[ACR⁺10, Theorem 1] because it does not require that $\sigma_+(\varphi) = O(n)$. For example, a formula that is biased near the root, but balanced at greater depths can have $\sigma_-(\varphi) = O(1)$ and $\sigma_+(\varphi) = \omega(n)$. By substituting the bound $\sigma_-(\varphi) = O(\sqrt{d})$ for a depth- d formula [ACR⁺10, Definition 3], a corollary of Theorem 8 is that a depth- d , size- n AND-OR formula can be evaluated using $O(\sqrt{nd})$ queries. This improves the depth-dependence from [ACR⁺10], and matches the dependence from an earlier version of that article [Amb07].

The essential reason that the Definition 1 balance condition can be weakened is that for the specific gates AND and OR, by writing out the optimal span programs explicitly we can prove that they satisfy stronger properties than are necessarily true for other functions.

2 Span Programs

2.1 Definitions

We briefly recall some definitions from [Rei09, Section 2]. Additionally, we define a span program complexity measure, the full witness size, that charges even for the “free” inputs. This quantity is important for developing quantum algorithms that are time efficient as well as query efficient.

For a natural number n , let $[n] = \{1, 2, \dots, n\}$. For a finite set X , let \mathbf{C}^X be the inner product space $\mathbf{C}^{|X|}$ with orthonormal basis $\{|x\rangle : x \in X\}$. For vector spaces V and W over \mathbf{C} , let $\mathcal{L}(V, W)$ be the set of linear transformations from V into W , and let $\mathcal{L}(V) = \mathcal{L}(V, V)$. For $A \in \mathcal{L}(V, W)$, $\|A\|$ is the operator norm of A . For a string $x \in \{0, 1\}^n$, let \bar{x} denote its bitwise complement.

Definition 3 ([HLŠ05, HLŠ07]). *For finite sets C, E and $\mathcal{D} \subseteq C^n$, let $f : \mathcal{D} \rightarrow E$. An adversary matrix for f is a real, symmetric matrix $\Gamma \in \mathcal{L}(\mathbf{C}^{\mathcal{D}})$ that satisfies $\langle x | \Gamma | y \rangle = 0$ whenever $f(x) \neq f(y)$.*

The general adversary bound for f , with costs $s \in [0, \infty)^n$, is

$$ADV_s^\pm(f) = \max_{\substack{\text{adversary matrices } \Gamma: \\ \forall j \in [n], \|\Gamma \circ \Delta_j\| \leq s_j}} \|\Gamma\|. \tag{2.1}$$

Here $\Gamma \circ \Delta_j$ denotes the entry-wise matrix product between Γ and $\Delta_j = \sum_{x,y: x_j \neq y_j} |x\rangle\langle y|$. The (nonnegative-weight) adversary bound for f , with costs s , is defined by the same maximization, except with Γ restricted to have nonnegative entries. In particular, $ADV_s^\pm(f) \geq \text{Adv}_s(f)$.

Letting $\vec{1} = (1, 1, \dots, 1)$, the adversary bound for f is $\text{Adv}(f) = \text{Adv}_{\vec{1}}(f)$ and the general adversary bound for f is $\text{Adv}^\pm(f) = \text{Adv}_{\vec{1}}^\pm(f)$. By [HLŠ07], $Q(f) = \Omega(\text{Adv}^\pm(f))$.

Definition 4 (Span program [KW93]). *A span program P consists of a natural number n , a finite-dimensional inner product space V over \mathbf{C} , a “target”*

vector $|t\rangle \in V$, disjoint sets I_{free} and $I_{j,b}$ for $j \in [n]$, $b \in \{0,1\}$, and “input vectors” $|v_i\rangle \in V$ for $i \in I_{\text{free}} \cup \bigcup_{j \in [n], b \in \{0,1\}} I_{j,b}$.

To P corresponds a function $f_P : \{0,1\}^n \rightarrow \{0,1\}$, defined on $x \in \{0,1\}^n$ by

$$f_P(x) = \begin{cases} 1 & \text{if } |t\rangle \in \text{Span}(\{|v_i\rangle : i \in I_{\text{free}} \cup \bigcup_{j \in [n]} I_{j,x_j}\}) \\ 0 & \text{otherwise} \end{cases} \quad (2.2)$$

Some additional notation is convenient. Fix a span program P . Let $I = I_{\text{free}} \cup \bigcup_{j \in [n], b \in \{0,1\}} I_{j,b}$. Let $A \in \mathcal{L}(\mathbf{C}^I, V)$ be given by $A = \sum_{i \in I} |v_i\rangle\langle i|$. For $x \in \{0,1\}^n$, let $I(x) = I_{\text{free}} \cup \bigcup_{j \in [n]} I_{j,x_j}$ and $\Pi(x) = \sum_{i \in I(x)} |i\rangle\langle i| \in \mathcal{L}(\mathbf{C}^I)$. Then $f_P(x) = 1$ if $|t\rangle \in \text{Range}(A\Pi(x))$. A vector $|w\rangle \in \mathbf{C}^I$ is said to be a witness for $f_P(x) = 1$ if $\Pi(x)|w\rangle = |w\rangle$ and $A|w\rangle = |t\rangle$. A vector $|w'\rangle \in V$ is said to be a witness for $f_P(x) = 0$ if $\langle t|w'\rangle = 1$ and $\Pi(x)A^\dagger|w'\rangle = 0$.

Definition 5 (Witness size). Consider a span program P , and a vector $s \in [0, \infty)^n$ of nonnegative “costs.” Let $S = \sum_{j \in [n], b \in \{0,1\}, i \in I_{j,b}} \sqrt{s_j} |i\rangle\langle i| \in \mathcal{L}(\mathbf{C}^I)$. For each input $x \in \{0,1\}^n$, define the witness size of P on x with costs s , $\text{wsz}_s(P, x)$, as follows:

$$\text{wsz}_s(P, x) = \begin{cases} \min_{|w\rangle: A\Pi(x)|w\rangle=|t\rangle} \|S|w\rangle\|^2 & \text{if } f_P(x) = 1 \\ \min_{\substack{|w'\rangle: \langle t|w'\rangle=1 \\ \Pi(x)A^\dagger|w'\rangle=0}} \|SA^\dagger|w'\rangle\|^2 & \text{if } f_P(x) = 0 \end{cases} \quad (2.3)$$

The witness size of P with costs s is

$$\text{wsz}_s(P) = \max_{x \in \{0,1\}^n} \text{wsz}_s(P, x) . \quad (2.4)$$

Define the full witness size $\text{fwsz}_s(P)$ by letting $S^f = S + \sum_{i \in I_{\text{free}}} |i\rangle\langle i|$ and

$$\text{fwsz}_s(P, x) = \begin{cases} \min_{|w\rangle: A\Pi(x)|w\rangle=|t\rangle} (1 + \|S^f|w\rangle\|^2) & \text{if } f_P(x) = 1 \\ \min_{\substack{|w'\rangle: \langle t|w'\rangle=1 \\ \Pi(x)A^\dagger|w'\rangle=0}} (\|w'\|^2 + \|SA^\dagger|w'\rangle\|^2) & \text{if } f_P(x) = 0 \end{cases} \quad (2.5)$$

$$\text{fwsz}_s(P) = \max_{x \in \{0,1\}^n} \text{fwsz}_s(P, x) . \quad (2.6)$$

When the subscript s is omitted, the costs are taken to be uniform, $s = \vec{1} = (1, 1, \dots, 1)$, e.g., $\text{fwsz}(P) = \text{fwsz}_{\vec{1}}(P)$. The witness size is defined in [RŠ08]. The full witness size is defined in [Rei09, Section 8], but is not named there. A *strict* span program has $I_{\text{free}} = \emptyset$, so $S^f = S$, and a *monotone* span program has $I_{j,0} = \emptyset$ for all j [Rei09, Definition 4.9].

2.2 Quantum Algorithm to Evaluate a Span Program Based on Its Full Witness Size

[Rei09, Theorem 9.3] gives a quantum query algorithm for evaluating span programs based on the full witness size. The algorithm is based on a quantum walk on a certain graph. Provided that the degree of the graph is not too large, it can actually be implemented efficiently.

Theorem 9 ([Rei09, Theorem 9.3]). *Let P be a span program. Then f_P can be evaluated using*

$$T = O(\text{fwsz}(P) \parallel \text{abs}(A_{G_P}) \parallel) \quad (2.7)$$

quantum queries, with error probability at most $1/3$. Moreover, if the maximum degree of a vertex in G_P is d , then the time complexity of the algorithm for evaluating f_P is at most a factor of $(\log d)(\log(T \log d))^{O(1)}$ worse, after classical preprocessing and assuming constant-time coherent access to the preprocessed string.

Proof. (sketch) The query complexity claim is actually slightly weaker than [Rei09, Theorem 9.3], which allows the target vector to be scaled downward by a factor of $\sqrt{\text{fwsz}(P)}$.

The time-complexity claim will follow from the proof of [Rei09, Theorem 9.3], in [Rei09, Prop. 9.4, Theorem 9.5]. The algorithm for evaluating $f_P(x)$ uses a discrete-time quantum walk on the graph $G_P(x)$. If the maximum degree of a vertex in G_P is d , then each coin reflection can be implemented using $O(\log d)$ single-qubit unitaries and queries to the preprocessed string [GR02, CNW10]. Finally, the $(\log(T \log d))^{O(1)}$ factor comes from applying the Solovay-Kitaev Theorem [KSV02] to compile the single-qubit unitaries into products of elementary gates, to precision $1/O(T \log d)$. \square

We remark that together with [Rei09, Theorem 3.1], Theorem 9 gives a way of transforming a one-sided-error quantum algorithm into a span program, and back into a quantum algorithm, such that the time complexity is nearly preserved, after preprocessing. This is only a weak equivalence, because aside from requiring preprocessing the algorithm from Theorem 9 also has two-sided error. To some degree, though, it complements the equivalence results for best span program witness size and bounded-error quantum *query* complexity [Rei09, Theorem 7.1, Theorem 9.2].

2.3 Direct-Sum Span Program Composition

Let us study the full witness size of the direct-sum composition of span programs. We begin by recalling the definition of direct-sum composition.

Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $S \subseteq [n]$. For $j \in [n]$, let m_j be a natural number, with $m_j = 1$ for $j \notin S$. For $j \in S$, let $f_j : \{0, 1\}^{m_j} \rightarrow \{0, 1\}$. Define $y : \{0, 1\}^{m_1} \times \cdots \times \{0, 1\}^{m_n} \rightarrow \{0, 1\}^n$ by

$$y(x)_j = \begin{cases} f_j(x_j) & \text{if } j \in S \\ x_j & \text{if } j \notin S \end{cases} \quad (2.8)$$

Define $g : \{0, 1\}^{m_1} \times \cdots \times \{0, 1\}^{m_n} \rightarrow \{0, 1\}$ by $g(x) = f(y(x))$. For example, if $S = [n] \setminus \{1\}$, then

$$g(x) = f(x_1, f_2(x_2), \dots, f_n(x_n)) \quad (2.9)$$

Given span programs for the individual functions f and f_j for $j \in S$, we will construct a span program for g . We remark that although we are here requiring that the inner functions f_j act on disjoint sets of bits, this assumption is not necessary for the definition. It simplifies the notation, though, for the cases $S \neq [n]$, and will suffice for our applications.

Let P be a span program computing $f_P = f$. Let P have inner product space V , target vector $|t\rangle$ and input vectors $|v_i\rangle$ indexed by I_{free} and I_{jc} for $j \in [n]$ and $c \in \{0, 1\}$.

For $j \in [n]$, let $s_j \in [0, \infty)^{m_j}$ be a vector of costs, and let $s \in [0, \infty)^{\sum m_j}$ be the concatenation of the vectors s_j . For $j \in S$, let P^{j0} and P^{j1} be span programs computing $f_{P^{j1}} = f_j : \{0, 1\}^{m_j} \rightarrow \{0, 1\}$ and $f_{P^{j0}} = \neg f_j$, with $r_j = \text{wsize}_{s_j}(P^{j0}) = \text{wsize}_{s_j}(P^{j1})$. For $c \in \{0, 1\}$, let P^{jc} have inner product space V^{jc} with target vector $|t^{jc}\rangle$ and input vectors indexed by I_{free}^{jc} and I_{kb}^{jc} for $k \in [m_j]$, $b \in \{0, 1\}$. For $j \notin S$, let $r_j = s_j$.

Let $I_S = \bigcup_{j \in S, c \in \{0, 1\}} I_{jc}$. Define $\varsigma : I_S \rightarrow [n] \times \{0, 1\}$ by $\varsigma(i) = (j, c)$ if $i \in I_{jc}$. The idea is that ς maps i to the input span program that must evaluate to 1 in order for $|v_i\rangle$ to be available in P .

There are several ways of composing the span programs P and P^{jc} to obtain a span program Q computing the composed function $f_Q = g$ with $\text{wsize}_s(Q) \leq \text{wsize}_r(P)$ [Rei09, Defs. 4.4, 4.5, 4.6]. We focus on direct-sum composition.

Definition 6 ([Rei09, Definition 4.5]). *The direct-sum-composed span program Q^\oplus is defined by:*

- The inner product space is $V^\oplus = V \oplus \bigoplus_{j \in S, c \in \{0, 1\}} (\mathbf{C}^{I_{jc}} \otimes V^{jc})$. Any vector in V^\oplus can be uniquely expressed as $|u\rangle_V + \sum_{i \in I_S} |i\rangle \otimes |u_i\rangle$, where $|u\rangle \in V$ and $|u_i\rangle \in V^{\varsigma(i)}$.
- The target vector is $|t^\oplus\rangle = |t\rangle_V$.
- The free input vectors are indexed by $I_{\text{free}}^\oplus = I_{\text{free}} \cup I_S \cup \bigcup_{j \in S, c \in \{0, 1\}} (I_{jc} \times I_{\text{free}}^{jc})$ with, for $i \in I_{\text{free}}^\oplus$,

$$|v_i^\oplus\rangle = \begin{cases} |v_i\rangle_V & \text{if } i \in I_{\text{free}} \\ |v_i\rangle_V - |i\rangle \otimes |t^{jc}\rangle & \text{if } i \in I_{jc} \text{ and } j \in S \\ |i'\rangle \otimes |v_{i''}\rangle & \text{if } i = (i', i'') \in I_{jc} \times I_{\text{free}}^{jc} \end{cases} \quad (2.10)$$

- The other input vectors are indexed by $I_{(jk)b}^\oplus$ for $j \in [n]$, $k \in [m_j]$, $b \in \{0, 1\}$. For $j \notin S$, $I_{(j1)b}^\oplus = I_{jb}$, with $|v_i^\oplus\rangle = |v_i\rangle_V$ for $i \in I_{(j1)b}^\oplus$. For $j \in S$, let $I_{(jk)b}^\oplus = \bigcup_{c \in \{0, 1\}} (I_{jc} \times I_{kb}^{jc})$. For $i \in I_{jc}$ and $i' \in I_{kb}^{jc}$, let

$$|v_{i'i'}^\oplus\rangle = |i\rangle \otimes |v_{i'}\rangle. \quad (2.11)$$

By [Rei09, Theorem 4.3], $f_{Q^\oplus} = g$ and $\text{wsize}_s(Q^\oplus) \leq \text{wsize}_r(P)$. (While that theorem is stated only for the case $S = [n]$, it is trivially extended to other $S \subset [n]$.) We give a bound on how quickly the full witness size can grow relative to the witness size:

Lemma 1. *Under the above conditions, for each input $x \in \{0, 1\}^{m_1} \times \dots \times \{0, 1\}^{m_n}$, with $y = y(x)$,*

- *If $g(x) = 1$, let $|w\rangle$ be a witness to $f_P(y) = 1$ such that*

$$\sum_{j \in [n], i \in I_{jy_j}} r_j |w_i|^2 = \text{wsize}_r(P, y).$$

Then

$$\frac{\text{fsize}_s(Q^\oplus, x)}{\text{wsize}_r(P, y)} \leq \sigma(y, |w\rangle) + \frac{1 + \sum_{i \in I_{\text{free}}} |w_i|^2}{\text{wsize}_r(P, y)}$$

$$\text{where } \sigma(y, |w\rangle) = \max_{\substack{j \in S: \\ \exists i \in I_{jy_j} \text{ with } \langle i | w \rangle \neq 0}} \frac{\text{fsize}_{s_j}(P^{jy_j})}{\text{wsize}_{s_j}(P^{jy_j})}. \quad (2.12)$$

- *If $g(x) = 0$, let $|w'\rangle$ be a witness to $f_P(y) = 0$ such that*

$$\sum_{j \in [n], i \in I_{j\bar{y}_j}} r_j |\langle w' | v_i \rangle|^2 = \text{wsize}_r(P, y).$$

Then

$$\frac{\text{fsize}_s(Q^\oplus, x)}{\text{wsize}_r(P, y)} \leq \sigma(\bar{y}, |w'\rangle) + \frac{\| |w'\rangle \|^2}{\text{wsize}_r(P, y)}$$

$$\text{where } \sigma(\bar{y}, |w'\rangle) = \max_{\substack{j \in S: \\ \exists i \in I_{j\bar{y}_j} \text{ with } \langle v_i | w' \rangle \neq 0}} \frac{\text{fsize}_{s_j}(P^{j\bar{y}_j})}{\text{wsize}_{s_j}(P^{j\bar{y}_j})}. \quad (2.13)$$

If $S = \emptyset$, then $\sigma(y, |w\rangle)$ and $\sigma(\bar{y}, |w'\rangle)$ should each be taken to be 1 in the above equations.

Proof. We follow the proof of [Rei09, Theorem 4.3], except keeping track of the full witness size. Note that if $S = \emptyset$, then Eqs. (2.12) and (2.13) are immediate by definition of $\text{fsize}_s(Q^\oplus, x)$.

Let $I(y)' = I(y) \setminus I_{\text{free}} = \bigcup_{j \in [n]} I_{jy_j}$.

In the first case, $g(x) = 1$, for $j \in S$ let $|w^{jy_j}\rangle \in \mathbf{C}^{I^{jy_j}}$ be a witness to $f_{P^{jy_j}}(x_j) = 1$ such that $\text{fsize}_s(P^{jy_j}, x_j) = 1 + \sum_{i \in I_{\text{free}}^{jy_j}} |w_i^{jy_j}|^2 + \sum_{k \in [m_j], i \in I_{k(x_j)_k}^{jy_j}} (s_j)_k |w_i^{jy_j}|^2$. As in [Rei09, Theorem 4.3], let $|w^\oplus\rangle \in \mathbf{C}^{I^\oplus(x)}$ be given by

$$w_i^\oplus = \begin{cases} w_i & \text{if } i \in I(y) \\ w_{i'} w_{i''}^{s(i')} & \text{if } i = (i', i'') \text{ with } i' \in I(y)' \cap I_S, i'' \in I^{s(i')}(x) \\ 0 & \text{otherwise} \end{cases} \quad (2.14)$$

Then $|w^\oplus\rangle$ is a witness for $f_{Q^\oplus}(x) = 1$, and we compute

$$\begin{aligned}
\text{fsize}_s(Q^\oplus, x) &\leq 1 + \sum_{i \in I_{\text{free}}^\oplus} |w_i^\oplus|^2 + \sum_{\substack{j \in [n], k \in [m_j], \\ i \in I_{(j,k)(x_j)_k}^\oplus}} (s_j)_k |w_i^\oplus|^2 \\
&= 1 + \sum_{i \in I_{\text{free}}} |w_i|^2 + \sum_{j \in [n] \setminus S, i \in I_{jx_j}} s_j |w_i|^2 \\
&\quad + \sum_{j \in S, i \in I_{jy_j}} |w_i|^2 \left(1 + \sum_{i' \in I_{\text{free}}^{jy_j}} |w_{i'}^{jy_j}|^2 \right. \\
&\quad \left. + \sum_{k \in [m_j], i' \in I_{k(x_j)_k}^{jy_j}} (s_j)_k |w_{i'}^{jy_j}|^2 \right) \\
&= 1 + \sum_{i \in I_{\text{free}}} |w_i|^2 + \sum_{j \in [n] \setminus S, i \in I_{jx_j}} s_j |w_i|^2 \\
&\quad + \sum_{j \in S, i \in I_{jy_j}} |w_i|^2 \text{fsize}_{s_j}(P^{jy_j}, x_j).
\end{aligned} \tag{2.15}$$

Equation (2.12) follows using the bound $\text{fsize}_{s_j}(P^{jy_j}, x_j) \leq \sigma(y, |w|)r_j$ for $j \in S$, and $s_j = r_j$ for $j \notin S$.

Next consider the case $g(x) = 0$. For $j \in S$, let $|u^{j\bar{y}_j}\rangle \in V^{j\bar{y}_j}$ be a witness for $f_{P^{j\bar{y}_j}}(x_j) = 0$ with $\text{fsize}_s(P^{j\bar{y}_j}, x_j) = \| |u^{j\bar{y}_j}\rangle \|^2 + \sum_{k \in [m_j], i \in I_{k(x_j)_k}^{j\bar{y}_j}} (s_j)_k |\langle v_i | u^{j\bar{y}_j} \rangle|^2$. As in [Rei09, Theorem 4.3], let

$$|u^\oplus\rangle = |w'\rangle_V + \sum_{i \in I_S \setminus I(y)} \langle v_i | w' \rangle |i\rangle \otimes |u^{s(i)}\rangle. \tag{2.16}$$

Then $|u^\oplus\rangle$ is a witness for $f_{Q^\oplus}(x) = 0$, and, moreover,

$$\begin{aligned}
\text{fsize}_s(Q^\oplus, x) &\leq \| |u^\oplus\rangle \|^2 + \sum_{j \in [n], k \in [m_j], i \in I_{(j,k)(x_j)_k}^\oplus} (s_j)_k |\langle v_i^\oplus | u^\oplus \rangle|^2 \\
&= \| |u^\oplus\rangle \|^2 + \sum_{\substack{j \in [n] \setminus S \\ i \in I_{j\bar{x}_j}}} s_j |\langle v_i^\oplus | u^\oplus \rangle|^2 \\
&\quad + \sum_{\substack{j \in S, k \in [m_j], \\ i \in I_{j\bar{y}_j}, i' \in I_{k(x_j)_k}^{j\bar{y}_j}}} (s_j)_k |\langle v_{i'}^\oplus | u^\oplus \rangle|^2
\end{aligned}$$

$$\begin{aligned}
 &= \| |w'\rangle \|^2 + \sum_{\substack{j \in [n] \setminus S \\ i \in I_{j\bar{x}_j}}} s_j |\langle v_i | w'\rangle|^2 & (2.17) \\
 &+ \sum_{j \in S, i \in I_{j\bar{y}_j}} |\langle v_i | w'\rangle|^2 \left(\| |u^{j\bar{y}_j}\rangle \|^2 \right. \\
 &+ \left. \sum_{k \in [m_j], i' \in I_{\frac{j\bar{y}_j}{k(x_j)_k}}} (s_j)_k |\langle v_{i'} | u^{j\bar{y}_j}\rangle|^2 \right) \\
 &= \| |w'\rangle \|^2 + \sum_{\substack{j \in [n] \setminus S \\ i \in I_{j\bar{x}_j}}} r_j |\langle v_i | w'\rangle|^2 \\
 &+ \sum_{j \in S, i \in I_{j\bar{y}_j}} |\langle v_i | w'\rangle|^2 \text{fwsz}_{s_j}(P^{j\bar{y}_j}, x_j) .
 \end{aligned}$$

Equation (2.13) follows using the bound $\text{fwsz}_{s_j}(P^{j\bar{y}_j}, x_j) \leq \sigma(\bar{y}, |w'\rangle) r_j$ for $j \in S$. \square

Lemma 1 is a key step in the formula-evaluation results in this article and [Rei11]. It is used to track the full witness size for span programs recursively composed in a direct-sum manner along a formula. The proof of Theorem 7 will require the lemma with the weaker bounds $\sigma(y, |w\rangle), \sigma(\bar{y}, |w'\rangle) \leq \max_{j \in S, c \in \{0,1\}} \text{fwsz}_{s_j}(P^{jc})/\text{wsz}_{s_j}(P^{jc})$. Theorem 8 will use only the slightly stronger bounds $\sigma(y, |w\rangle) \leq \max_{j \in S} \text{fwsz}_{s_j}(P^{jy_j})/\text{wsz}_{s_j}(P^{jy_j}), \sigma(\bar{y}, |w'\rangle) \leq \max_{j \in S} \text{fwsz}_{s_j}(P^{j\bar{y}_j})/\text{wsz}_{s_j}(P^{j\bar{y}_j})$. However, the proof of [Rei11, Theorem 1.1] will require the bounds of Eqs. (2.12) and (2.13).

3 Evaluation of Almost-Balanced Formulas

In this section, we will apply the span program framework from [Rei09] to prove Theorem 7. Our algorithm will be given by applying Theorem 9 to a certain span program. Before beginning the proof, though, we will give two necessary lemmas.

Consider a span program P with corresponding weighted graph G_P , from [Rei09, Definition 8.2]. We will need a bound on the operator norm of $\text{abs}(A_{G_{P_v}})$, the entry-wise absolute value of the weighted adjacency matrix $A_{G_{P_v}}$. If P is *canonical* [Rei09, Definition 5.1], then we can indeed obtain such a bound in terms of the witness size of P :

Lemma 2. *Let $s \in (0, \infty)^k$, and let P be a canonical span program computing a function $f : \{0, 1\}^k \rightarrow \{0, 1\}$ with input vectors indexed by the set I . Assume that for each $x \in \{0, 1\}^k$ with $f(x) = 0$, an optimal witness to $f_P(x) = 0$ is $|x\rangle$ itself. Then*

$$\| \text{abs}(A_{G_P}) \| \leq 2^k \left(1 + \frac{\text{wsz}_s(P)}{\min_{j \in [k]} s_j} \right) + |I| . \tag{3.1}$$

Proof. Recall from [Rei09, Definition 5.1], that P being in canonical form implies that its target vector is $|t\rangle = \sum_{x:f(x)=0} |x\rangle$, and that the matrix A whose columns are the input vectors of P can be expressed as

$$A = \sum_{i \in I} |v_i\rangle\langle i| = \sum_{j \in [k], x: f(x)=0} |x\rangle\langle j, \bar{x}_j| \otimes \langle v_{xj}| . \quad (3.2)$$

By assumption, for each $x \in f^{-1}(0)$,

$$\sum_{j \in [k]} s_j \| |v_{xj}\rangle \|^2 = \text{wsize}_s(P, x) \leq \text{wsize}_s(P) . \quad (3.3)$$

In particular, letting $\sigma = \min_{j \in [k]} s_j > 0$, we can bound

$$\begin{aligned} \sum_{j \in [k]} \| |v_{xj}\rangle \|^2 &\leq \frac{1}{\sigma} \sum_{j \in [k]} s_j \| |v_{xj}\rangle \|^2 \\ &\leq \frac{\text{wsize}_s(P)}{\sigma} . \end{aligned} \quad (3.4)$$

The rest of the argument follows from the definition of the weighted adjacency matrix A_{G_P} . From [Rei09, Definition 8.1, Prop. 8.8], $\| \text{abs}(A_{G_P}) \| \leq \| \text{abs}(B_{G_P}) \|^2$, where B_{G_P} is the *biadjacency* matrix corresponding to P ,

$$B_{G_P} = \begin{pmatrix} |t\rangle & A \\ 0 & \mathbf{1} \end{pmatrix} , \quad (3.5)$$

and $\mathbf{1}$ is an $|I| \times |I|$ identity matrix. Now bound $\| \text{abs}(B_{G_P}) \|$ by its Frobenius norm:

$$\begin{aligned} \| \text{abs}(A_{G_P}) \| &\leq \| \text{abs}(B_{G_P}) \|^2 \\ &\leq \| \text{abs}(B_{G_P}) \|_F^2 \\ &= \| |t\rangle \|^2 + \sum_{\substack{x: f(x)=0, \\ j \in [k]}} \| |v_{xj}\rangle \|^2 + |I| \\ &\leq 2^k + 2^k \max_{x: f(x)=0} \sum_{j \in [k]} \| |v_{xj}\rangle \|^2 + |I| . \end{aligned} \quad (3.6)$$

Equation (3.1) follows by substituting in Eq. (3.4). \square

An important quantity in the proof of Theorem 7 will be $\sigma_{-}(\varphi)$, from Definition 2. For an almost-balanced formula φ , $\sigma_{-}(\varphi) = O(1)$.

Lemma 3. *Consider a β -balanced formula φ over a gate set \mathcal{S} in which every gate depends on at least two input bits. Then for every vertex v , with children c_1, c_2, \dots, c_k ,*

$$\frac{\text{Adv}^{\pm}(\varphi_v)}{\max_j \text{Adv}^{\pm}(\varphi_{c_j})} \geq \sqrt{1 + \frac{1}{\beta^2}} . \quad (3.7)$$

In particular,

$$\sigma_{-}(\varphi) \leq (2 + \sqrt{2})\beta^2 . \quad (3.8)$$

Proof. Consider a vertex v with corresponding gate $g = g_v : \{0, 1\}^k \rightarrow \{0, 1\}$. By Theorem 2, $\text{Adv}^\pm(\varphi_v) = \text{Adv}^\pm_s(g)$, where $s_j = \text{Adv}^\pm(\varphi_{c_j})$. It is immediate from the definitions that $\text{ADV}_s^\pm(g) \geq \text{Adv}_s(g)$. We will show that $\text{Adv}_s(g) \geq \sqrt{1 + 1/\beta^2}(\max_j s_j)$, using that $\max_j s_j / \min_j s_j \leq \beta$.

Use the weighted minimax formulation of the adversary bound from [HLS07, Theorem 18]:

$$\text{Adv}_s(g) = \min_P \max_{\substack{x, y \in \{0, 1\}^k \\ g(x) \neq g(y)}} \frac{1}{\sum_{j: x_j \neq y_j} \sqrt{p_x(j)p_y(j)/s_j}}, \tag{3.9}$$

where the minimization is over all choices of probability distributions p_x over $[k]$ for $x \in \{0, 1\}^k$.

Since the adversary bound is monotone increasing in each weight, the worst case is when all but one of the weights are equal to $\max_j s_j / \beta$. Since for a scalar c , $\text{Adv}_{cs}(g) = c\text{Adv}_s(g)$, we may scale so that one weight is β and all other weights are 1. Assume that the first weight is $s_1 = \beta$; the other $k - 1$ cases, $s_2 = \beta$ and so on, are symmetrical. Assume also that g depends on the first bit; otherwise $\text{Adv}_s^\pm(g)$ will not depend on s_1 so one of the other cases will be worse. Therefore, there exist inputs $x, y \in \{0, 1\}^k$ that differ only on the first bit, but for which $g(x) \neq g(y)$.

Since the function g depends on at least two input bits, there also exists a third input $z \in \{0, 1\}^k$ with $x_1 = z_1$ but $g(z) = g(y) \neq g(x)$. Indeed, if $g(z) = g(x)$ for every z with $z_1 = x_1$, and if $g(z) = g(y)$ for every z with $z_1 = y_1$, then g depends only on the first bit.

By Eq. (3.9),

$$\text{ADV}_s^\pm(g) \geq \min_{p_x, p_y, p_z} \max \left\{ \frac{1}{\sqrt{p_x(1)p_y(1)/s_1}}, \frac{1}{\sum_{\substack{j \geq 2 \\ x_j \neq z_j}} \sqrt{p_x(j)p_z(j)/s_j}} \right\} \tag{3.10}$$

where the minimization is over only the three probability distributions p_x, p_y and p_z . In the above expression, we may clearly take $p_y(1) = 1$ and $p_y(j) = 0$ for $j \geq 2$. We may also use the Cauchy-Schwarz inequality to bound the second term above, and finally substitute $s_1 = \beta, s_j = 1$ for $j \geq 2$ to obtain,

$$\text{ADV}_s^\pm(g) \geq \min_{p_x} \max \left\{ \frac{\beta}{\sqrt{p_x(1)}}, \frac{1}{\sqrt{\sum_{j \geq 2} p_x(j)}} \right\}. \tag{3.11}$$

The optimum is achieved for $p_x(1) = \beta^2 / (1 + \beta^2)$, so $\text{Adv}_s^\pm(g) \geq \sqrt{1 + \beta^2}$, as claimed.

To derive Eq. (3.8), note that $\beta \geq 1$ necessarily. Then the sum $\sigma_-(\varphi)$ is dominated by the geometric series

$$\sum_{k=0}^{\infty} \left(1 + \frac{1}{\beta^2}\right)^{-k/2}, \tag{3.12}$$

which is at most $(2 + \sqrt{2})\beta^2$, with equality at $\beta = 1$. □

Note that the 1-balanced formulas over $\mathcal{S} = \{\text{OR}_2\}$ satisfy the inequality (3.7) with equality and come arbitrarily close to saturating the inequality (3.8).

With Lemmas 2 and 3 in hand, we are ready to prove Theorem 7.

Proof. (of Theorem 7) First of all, we may assume without loss of generality that every gate in \mathcal{S} depends on at least two input bits. Indeed, if a gate $g : \{0, 1\}^k \rightarrow \{0, 1\}$ depends on no input bits, i.e., is the constant 0 or constant 1 function, then g can be eliminated from any formula over \mathcal{S} without changing the adversary balance condition, since $\text{ADV}_s^\pm(g) = 0$ for all cost vectors $s \in [0, \infty)^k$. If a gate $g : \{0, 1\}^k \rightarrow \{0, 1\}$ depends only on one input bit, say the first bit, then $\text{ADV}_s^\pm(g) = s_1$ for all cost vectors s , and therefore similarly g can be eliminated without affecting the adversary balance condition.

Consider φ an n -variable, β -balanced, read-once formula over the finite gate set \mathcal{S} . Let r be the root of φ . We begin by recursively constructing a span program P_φ that computes φ and has witness size $\text{wsiz}(P_\varphi) = \text{Adv}^\pm(\varphi)$. P_φ is constructed using direct-sum composition of span programs for each node in φ . (Direct-sum composition is also the composition method used in [RŠ08].)

The construction works recursively, starting at the leaves of φ and moving toward the root. Consider an internal vertex v , with children c_1, \dots, c_k . Let $\alpha_j = \text{Adv}^\pm(\varphi_{c_j})$, where φ_{c_j} is the subformula of φ rooted at c_j (Definition 1). In particular, if c_j is a leaf, then $\alpha_j = 1$. Assume that for $j \in [k]$ we have inductively constructed span programs $P_{\varphi_{c_j}}$ and $P_{\varphi_{c_j}}^\dagger$ computing φ_{c_j} and $\neg\varphi_{c_j}$, respectively, with $\text{wsiz}(P_{\varphi_{c_j}}) = \text{wsiz}(P_{\varphi_{c_j}}^\dagger) = \alpha_j$. Apply [Rei09, Theorem 6.1], a generalization of Theorem 4, twice to obtain span programs P_v and P_v^\dagger computing $f_{P_v} = g_v$ and $f_{P_v^\dagger} = \neg g_v$, with $\text{wsiz}_\alpha(P_v) = \text{wsiz}_\alpha(P_v^\dagger) = \text{ADV}_\alpha^\pm(g_v) = \text{Adv}^\pm(\varphi_v)$.

Then let P_{φ_v} and $P_{\varphi_v}^\dagger$ be the direct-sum-composed span programs of P_v and P_v^\dagger , respectively, with the span programs $P_{\varphi_{c_j}}$, $P_{\varphi_{c_j}}^\dagger$ according to the formula φ . By definition of direct-sum composition, the graph $G_{P_{\varphi_v}}$ is built by replacing the input edges of G_{P_v} with the graphs $G_{P_{\varphi_{c_j}}}$ or $G_{P_{\varphi_{c_j}}^\dagger}$; and similarly for $G_{P_{\varphi_v}^\dagger}$. Some examples are given in [Rei09, Appendix B] and in [RŠ08]. By [Rei09, Theorem 4.3], P_{φ_v} (resp. $P_{\varphi_v}^\dagger$) computes φ_v ($\neg\varphi_v$) with $\text{wsiz}(P_{\varphi_v}) = \text{wsiz}(P_{\varphi_v}^\dagger) = \text{Adv}^\pm(\varphi_v)$.

Let $P_\varphi = P_{\varphi_r}$. We wish to apply Theorem 9 to P_φ to obtain a quantum algorithm, but to do so will need some more properties of the span programs P_v and P_v^\dagger . Recall from [Rei09, Theorem 5.2] that each P_v may be assumed to be in canonical form, satisfying in particular that for any input $y \in \{0, 1\}^k$ with $g_v(y) = 0$ an optimal witness is $|y\rangle \in \mathbf{C}^{g_v^{-1}(0)}$ itself. Therefore, Lemma 2 applies, and we obtain

$$\| \text{abs}(A_{G_{P_v}}) \| = 2^k \left(1 + \frac{\text{wsiz}_\alpha(P_v)}{\min_j \alpha_j} \right) + |I|, \quad (3.13)$$

where $|I|$ is the number of input vectors in P_v . Now use

$$\begin{aligned} \frac{\text{wsize}_\alpha(P_v)}{\min_j \alpha_j} &= \frac{\max_j \alpha_j \text{Adv}^\pm_\alpha(g_v)}{\min_j \alpha_j \max_j \alpha_j} \\ &\leq \beta k, \end{aligned} \tag{3.14}$$

where we have applied Eq. (1.10) and also $\text{ADV}^\pm_\alpha(g_v)/\max_j \alpha_j \leq \text{Adv}^\pm(g_v) \leq k$. Additionally, by [Rei09, Lemma 6.6], we may assume that $|I| \leq 2k^2 2^k$. Thus

$$\|\text{abs}(A_{G_{P_v}})\| = \beta 2^{O(k)}. \tag{3.15}$$

By repeating this argument for the negated function $\neg g_v$ computed by a dual span program P_v^\dagger ([Rei09, Lemma 4.1]), we also have $\|\text{abs}(A_{G_{P_v^\dagger}})\| = \beta 2^{O(k)}$.

A consequence is that

$$\|\text{abs}(A_{G_{P_\varphi}})\| = \beta 2^{O(k_{\max})} \tag{3.16}$$

where k_{\max} is the maximum fan-in of any gate used in φ . Indeed, G_{P_φ} is built by “plugging together” the graphs G_{P_v} and $G_{P_v^\dagger}$ for the different vertices v . Split the graph G_{P_φ} into two pieces, G_0 and G_1 , comprising those subgraphs G_{P_v} and $G_{P_v^\dagger}$ for which the distance of v from r is even or odd, respectively. Then $\|\text{abs}(A_{G_{P_\varphi}})\| \leq \|\text{abs}(A_{G_0})\| + \|\text{abs}(A_{G_1})\|$. Since each G_b is the disconnected union of graphs G_{P_v} and $G_{P_v^\dagger}$, $\|\text{abs}(A_{G_b})\| \leq \max_v \max\{\|\text{abs}(A_{G_{P_v}})\|, \|\text{abs}(A_{G_{P_v^\dagger}})\|\}$.

Let us bound the full witness size of P_φ .

Lemma 4. *Let v be a vertex of φ . Then*

$$\max\{\text{fwsz}(P_{\varphi_v}), \text{fwsz}(P_{\varphi_v}^\dagger)\} \leq \sigma_-(v) \text{Adv}^\pm(\varphi_v). \tag{3.17}$$

Proof. The proof is by induction in the maximum distance from v to a leaf. The base case, that all of v ’s inputs are themselves leaves is by definition of P_v and P_v^\dagger , since then $\sigma_-(v) = 1 + 1/\text{Adv}^\pm(g_v)$.

Let v have children c_1, \dots, c_k . By Lemma 1 with $s = \overline{1}$ and $S = \{j \in [k] : c_j \text{ is not a leaf}\}$,

$$\frac{\text{fwsz}(P_{\varphi_v})}{\text{Adv}^\pm(\varphi_v)} \leq \frac{1}{\text{Adv}^\pm(\varphi_v)} + \max_{j \in S} \max \left\{ \frac{\text{fwsz}(P_{\varphi_{c_j}})}{\text{Adv}^\pm(\varphi_{c_j})}, \frac{\text{fwsz}(P_{\varphi_{c_j}}^\dagger)}{\text{Adv}^\pm(\varphi_{c_j})} \right\}. \tag{3.18}$$

In the case $\varphi_v(x) = 1$, this follows since P_v is strict, so in Eq. (2.12) the sum over I_{free} is zero. In the case $\varphi_v(x) = 0$, this follows since P_v is in canonical form, so in Eq. (2.13), $\|w'\|^2 = 1$.

Now by induction, the right-hand side is at most $\text{Adv}^\pm(\varphi_v)^{-1} + \max_{j \in S} \sigma_-(\varphi_{c_j}) = \sigma_-(v)$. \square

In particular, applying Lemma 4 for the case $v = r$, we find

$$\text{fwsz}(P_\varphi) \leq \sigma_-(\varphi) \text{Adv}^\pm(\varphi) = O(\beta^2 \text{Adv}^\pm(\varphi)) \tag{3.19}$$

since $\sigma_-(\varphi) = O(\beta^2)$ by Lemma 3. Combining Eqs. (3.16) and (3.19) gives

$$\text{fsize}(P_\varphi) \|\text{abs}(A_{G_{P_\varphi}})\| = \beta^3 2^{O(k_{\max})} \text{Adv}^\pm(\varphi) . \quad (3.20)$$

This is $O(\text{Adv}^\pm(\varphi))$; since the gate set \mathcal{S} is fixed and finite, $k_{\max} = O(1)$. Theorem 7 now follows from Theorem 9. \square

Note that the lost constant in the theorem grows cubically in the balance parameter β and exponentially in the maximum fan-in k_{\max} of a gate in \mathcal{S} . It is conceivable that this exponential dependence can be improved.

For future reference, we state separately the bound used above to derive Eq. (3.16).

Lemma 5. *If P_φ is the direct-sum composition along a formula φ of span programs P_v and P_v^\dagger , then*

$$\|\text{abs}(A_{G_P})\| \leq 2 \max_{v \in \varphi} \max\{\|\text{abs}(A_{G_{P_v}})\|, \|\text{abs}(A_{G_{P_v^\dagger}})\|\} . \quad (3.21)$$

If the span programs P_v are monotone, then $\|\text{abs}(A_{G_P})\| \leq 2 \max_v \|\text{abs}(A_{G_{P_v}})\|$.

The claim for monotone span programs follows because then the dual span programs P_v^\dagger are not used in P_φ .

4 Evaluation of Approximately Balanced AND-OR Formulas

The proof of Theorem 8 will again be a consequence of Lemma 1 and Theorem 9.

We will use the following strict, monotone span programs for fan-in-two AND and OR gates:

Definition 7. *For $s_1, s_2 > 0$, define span programs $P_{\text{AND}}(s_1, s_2)$ and $P_{\text{OR}}(s_1, s_2)$ computing AND_2 and OR_2 , $\{0, 1\}^2 \rightarrow \{0, 1\}$, respectively, by*

$$P_{\text{AND}}(s_1, s_2) : |t\rangle = \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix}, \quad |v_1\rangle = \begin{pmatrix} \beta_1 \\ 0 \end{pmatrix}, \quad |v_2\rangle = \begin{pmatrix} 0 \\ \beta_2 \end{pmatrix} \quad (4.1)$$

$$P_{\text{OR}}(s_1, s_2) : |t\rangle = \delta, \quad |v_1\rangle = \epsilon_1, \quad |v_2\rangle = \epsilon_2 \quad (4.2)$$

Both span programs have $I_{1,1} = \{1\}$, $I_{2,1} = \{2\}$ and $I_{\text{free}} = I_{1,0} = I_{2,0} = \emptyset$. Here the parameters $\alpha_j, \beta_j, \delta, \epsilon_j$, for $j \in [2]$, are given by

$$\alpha_j = (s_j/s_p)^{1/4} \quad \beta_j = 1 \quad (4.3)$$

$$\delta = 1 \quad \epsilon_j = (s_j/s_p)^{1/4} , \quad (4.4)$$

where $s_p = s_1 + s_2$. Let $\alpha = \sqrt{\alpha_1^2 + \alpha_2^2}$ and $\epsilon = \sqrt{\epsilon_1^2 + \epsilon_2^2}$.

Note that $\alpha, \epsilon \in (1, 2^{1/4}]$. They are largest when $s_1 = s_2$.

Claim. The span programs $P_{\text{AND}}(s_1, s_2)$ and $P_{\text{OR}}(s_1, s_2)$ satisfy:

$$\begin{aligned} \text{wsize}_{(\sqrt{s_1}, \sqrt{s_2})}(P_{\text{AND}}, x) &= \begin{cases} \sqrt{s_p} & \text{if } x \in \{11, 10, 01\} \\ \frac{\sqrt{s_p}}{2} & \text{if } x = 00 \end{cases} \\ \text{wsize}_{(\sqrt{s_1}, \sqrt{s_2})}(P_{\text{OR}}, x) &= \begin{cases} \sqrt{s_p} & \text{if } x \in \{00, 10, 01\} \\ \frac{\sqrt{s_p}}{2} & \text{if } x = 11 \end{cases} \end{aligned} \quad (4.5)$$

Proof. These are calculations using Definition 5 for the witness size. Letting $\sigma = (\sqrt{s_1}, \sqrt{s_2})$, $Q = P_{\text{AND}}(s_1, s_2)$ and $R = P_{\text{OR}}(s_1, s_2)$, we have

$$\text{wsize}_{\sigma}(Q, 11) = \left(\frac{\alpha_1}{\beta_1}\right)^2 \sqrt{s_1} + \left(\frac{\alpha_2}{\beta_2}\right)^2 \sqrt{s_2} = \sqrt{s_p} \quad (4.6)$$

$$\text{wsize}_{\sigma}(Q, 10) = \left(\frac{\beta_2}{\alpha_2}\right)^2 \sqrt{s_2} = \sqrt{s_p} \quad (4.7)$$

$$\text{wsize}_{\sigma}(Q, 00) = \left(\left(\frac{\alpha_1}{\beta_1}\right)^2 \frac{1}{\sqrt{s_1}} + \left(\frac{\alpha_2}{\beta_2}\right)^2 \frac{1}{\sqrt{s_2}} \right)^{-1} = \frac{\sqrt{s_p}}{2} \quad (4.8)$$

$$\text{wsize}_{\sigma}(Q, 01) = \left(\frac{\beta_1}{\alpha_1}\right)^2 \sqrt{s_1} = \sqrt{s_p} \quad (4.9)$$

and

$$\text{wsize}_{\sigma}(R, 11) = \delta^2 \left(\frac{\epsilon_1^2}{\sqrt{s_1}} + \frac{\epsilon_2^2}{\sqrt{s_2}} \right)^{-1} = \frac{\sqrt{s_p}}{2} \quad (4.10)$$

$$\text{wsize}_{\sigma}(R, 10) = \left(\frac{\delta}{\epsilon_1}\right)^2 \sqrt{s_1} = \sqrt{s_p} \quad (4.11)$$

$$\text{wsize}_{\sigma}(R, 00) = \left(\frac{\epsilon_1}{\delta}\right)^2 \sqrt{s_1} + \left(\frac{\epsilon_2}{\delta}\right)^2 \sqrt{s_2} = \sqrt{s_p} \quad (4.12)$$

$$\text{wsize}_{\sigma}(R, 01) = \left(\frac{\delta}{\epsilon_2}\right)^2 \sqrt{s_2} = \sqrt{s_p} . \quad (4.13)$$

It is not a coincidence that $\text{wsize}_{\sigma}(Q, x) = \text{wsize}_{\sigma}(R, \bar{x})$ for all $x \in \{0, 1\}^2$. This can be seen as a consequence of De Morgan's laws and span program duality—see [Rei09, Lemma 4.1]. \square

Proof. (of Theorem 8) Let φ be an AND-OR formula of size n , i.e., on n input bits.

First expand out the formula so that every AND gate and every OR gate has fan-in two. This expansion can be carried out without increasing $\sigma_{-}(\varphi)$ by more than a factor of 10:

Lemma 6 ([ACR⁺10, Lemma 8]) *For any AND-OR formula φ , one can efficiently construct an equivalent AND-OR formula φ' of the same size, such that all gates in φ' have fan-in at most two, and $\sigma_{-}(\varphi') = O(\sigma_{-}(\varphi))$.*

Therefore we may assume that φ is a formula over fan-in-two AND and OR gates.

Now use direct-sum composition to compose the AND and OR gates according to the formula φ , as in the proof of Theorem 7. Since the span programs for AND and OR are monotone, direct-sum composition does not make use of dual span programs computing NAND or NOR. Therefore there is no need to specify these span programs. At a vertex v , set the weights s_1 and s_2 to equal the sizes of v 's two input subformulas. Let P_v be the span program used at vertex v , P_{φ_v} be the span program thus constructed for the subformula φ_v , and P_φ be the span program constructed computing φ . With this choice of weights, it follows from Claim 4 and [Rei09, Theorem 4.3] that $\text{wsize}(P_{\varphi_v}) = \text{Adv}^\pm(\varphi_v) = \text{Adv}(\varphi_v)$.

Notice that for all $s_1, s_2 \in [0, \infty)$, $\|\text{abs}(A_{G_{P_{\text{AND}}}(s_1, s_2)}})\| = O(1)$ and $\|\text{abs}(A_{G_{P_\varphi}})\| = O(1)$. Therefore, by Lemma 5, we obtain that $\|\text{abs}(A_{G_{P_\varphi}})\| = O(1)$.

Thus to apply Theorem 9 we need only bound $\text{fsize}(P_\varphi)$. Lemma 4 does not apply, because for $P_{\text{AND}}(s_1, s_2)$, an optimal witness $|w'\rangle$ to $f_{P_{\text{AND}}}(x) = 0$ might have $\| |w'\rangle \|^2 > 1$, as each $\alpha_j < 1$. (Lemma 4 would apply had we set the parameters to be $\alpha_1 = \alpha_2 = 1$, $\beta_j = (s_p/s_j)^{1/4}$, but then $\|A_{G_{P_{\text{AND}}}}\|$ would not necessarily be $O(1)$.) However, analogous to Lemma 4, we will show:

Lemma 7. *Let v be a vertex of φ . Then*

$$\text{fsize}(P_{\varphi_v}, x) \leq \begin{cases} \sigma_-(v)\text{Adv}(\varphi_v) & \text{if } \varphi_v(x) = 1 \\ 2\sigma_-(v)\text{Adv}(\varphi_v) - 1 & \text{if } \varphi_v(x) = 0 \end{cases} \quad (4.14)$$

Proof. The proof is by induction in the maximum distance from v to a leaf. The base case, that v 's two inputs are themselves leaves is by definition of P_v , since then $\sigma_-(v) = 1 + 1/\sqrt{2}$.

Let v have children c_1 and c_2 . We will use Lemma 1 with $s = \vec{1}$, $S = \{j \in [2] : c_j \text{ is not a leaf}\}$.

If $\varphi_v(x) = 1$, then since P_v is a strict span program, i.e., $I_{\text{free}} = \emptyset$, Eq. (2.12) gives

$$\frac{\text{fsize}(P_{\varphi_v}, x)}{\text{Adv}(\varphi_v)} \leq \frac{1}{\text{Adv}(\varphi_v)} + \max_{j \in S} \frac{\text{fsize}(P_{\varphi_{c_j}})}{\text{Adv}(\varphi_{c_j})}. \quad (4.15)$$

By induction, the right-hand side is at most $1/\text{Adv}(\varphi_v) + \max_j \sigma_-(c_j) = \sigma_-(v)$.

If $\varphi_v(x) = 0$ and g_v is an OR gate, then the unique witness $|w'\rangle$ for P_v has $\| |w'\rangle \| = 1$, from Definition 7. From Eq. (2.13) and the induction hypothesis,

$$\begin{aligned} \frac{\text{fsize}(P_{\varphi_v}, x)}{\text{Adv}^\pm(\varphi_v)} &\leq \frac{1}{\text{Adv}(\varphi_v)} + \max_{j \in S} \left(2\sigma_-(c_j) - \frac{1}{\text{Adv}(\varphi_{c_j})} \right) \\ &< 2\sigma_-(v) - \frac{1}{\text{Adv}(\varphi_v)}, \end{aligned} \quad (4.16)$$

as claimed.

Therefore assume that $\varphi_v(x) = 0$ and g_v is an AND gate. Let s_1 and s_2 be the sizes of the two input subformulas to v , $s_p = s_1 + s_2 = \text{Adv}(\varphi_v)^2$, and assume without loss of generality that $\varphi_{c_1}(x) = 0$. If $\varphi_{c_2}(x) = 0$ as well, then assume without loss of generality that $2\sigma_-(c_1) - \frac{1}{\sqrt{s_1}} \geq 2\sigma_-(c_2) - \frac{1}{\sqrt{s_2}}$, so $\sigma(\bar{y}) \leq 2\sigma_-(c_1) - \frac{1}{\sqrt{s_1}}$. Then the witness $|w'\rangle$ may be taken to be $|w'\rangle = (1/\alpha_1, 0) = ((s_p/s_1)^{1/4}, 0)$. From Eq. (2.13),

$$\begin{aligned} \frac{\text{fwsiz}(P_{\varphi_v}, x)}{\text{Adv}^\pm(\varphi_v)} &\leq \frac{\sqrt{s_p/s_1}}{\text{Adv}^\pm(\varphi_v)} + \sigma(\bar{y}) \\ &\leq \frac{1}{\sqrt{s_1}} + \left(2\sigma_-(c_1) - \frac{1}{\sqrt{s_1}}\right) \\ &< 2\sigma_-(v) - \frac{1}{\sqrt{s_p}}, \end{aligned} \tag{4.17}$$

as claimed. \square

In particular, applying Lemma 7 for the case $v = r$, we find

$$\text{fwsiz}(P_\varphi) \leq 2\sigma_-(\varphi)\text{Adv}(\varphi) = 2\sigma_-(\varphi)\sqrt{n}. \tag{4.18}$$

Theorem 8 now follows from Theorem 9. \square

5 Open Problems

In order to begin to relax the balance condition for general formulas, it seems that we need a better understanding of the canonical span programs. For example, can the norm bound Lemma 2 be improved?

Although the two-sided bounded-error quantum query complexity of evaluating formulas is beginning to be understood, the zero-error quantum query complexity [BCWZ99] appears to be more complicated. For example, the exact and zero-error quantum query complexities for OR_n are both n [BBC⁺01]. On the other hand, Ambainis et al. [ACGT10] use the [ACR⁺10] algorithm as a subroutine in the construction of a self-certifying, zero-error quantum algorithm that makes $O(\sqrt{n} \log^2 n)$ queries to evaluate the balanced binary AND-OR formula. It is not known how to relax the balance requirement or extend the gate set.

Can we develop further methods for constructing span programs with small full witness size, norm and maximum degree? A companion paper [Rei11] studies reduced tensor-product span program composition in order to complement the direct-sum composition that we have used here.

The case of formulas over non-boolean gates may be more complicated [Rei09], but is still intriguing.

Acknowledgements. I thank Andrew Landahl and Robert Špalek for helpful discussions. Research supported by NSERC and ARO-DTO.

References

- [ACGT10] Ambainis, A., Childs, A.M., Le Gall, F., Tani, S.: The quantum query complexity of certification. *Quantum Inf. Comput.* **10**, 181–188 (2010) (arXiv:0903.1291[quant-ph])
- [ACR⁺10] Ambainis, A., Childs, A.M., Reichardt, B.W., Špalek, R., Zhang, S.: Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM J. Comput.* **39**(6), 2513–2530 (2010). (Earlier version in FOCS'07)
- [Amb02] Ambainis, A.: Quantum lower bounds by quantum arguments. *J. Comput. Syst. Sci.* **64**, 750–767 (2002). (arXiv:quant-ph/0002066. Earlier version in STOC'00)
- [Amb05] Ambainis, A.: Polynomial degree and lower bounds in quantum complexity: collision and element distinctness with small range. *Theory Comput.* **1**, 37–46 (2005). (arXiv:quant-ph/0305179)
- [Amb06] Ambainis, A.: Polynomial degree vs. quantum query complexity. *J. Comput. Syst. Sci.* **72**(2), 220–238 (2006). (arXiv:quant-ph/0305028. Earlier version in FOCS'03)
- [Amb07] Ambainis, A.: A nearly optimal discrete query quantum algorithm for evaluating NAND formulas (2007). (arXiv:0704.3628[quant-ph])
- [AS04] Aaronson, S., Shi, Y.: Quantum lower bounds for the collision and the element distinctness problem. *J. ACM* **51**(4), 595–605 (2004)
- [BB94] Bonet, M.L., Buss, S.R.: Size-depth tradeoffs for Boolean. *Inf. Process. Lett.* **49**(3), 151–155 (1994)
- [BBBV97] Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.: Strengths and weaknesses of quantum computing. *SIAM J. Comput.* **26**(5), 1510–1523 (1997). (arXiv:quant-ph/9701001)
- [BBC⁺01] Beals, R., Buhrman, H., Cleve, R., Mosca, M., de Wolf, R.: Quantum lower bounds by polynomials. *J. ACM* **48**(4), 778–797 (2001). (arXiv:quant-ph/9802049. Earlier version in FOCS'98)
- [BCE91] Bshouty, N.H., Cleve, R., Eberly, W.: Size-depth tradeoffs for algebraic formulae. In: *Proceedings of the 32nd IEEE FOCS*, pp. 334–341 (1991)
- [BCW98] Buhrman, H., Cleve, R., Wigderson, A.: Quantum vs. classical communication and computation. In: *Proceedings of the 30th ACM STOC*, pp. 63–68 (1998) (arXiv:quant-ph/9802040)
- [BCWZ99] Buhrman, H., Cleve, R., de Wolf, R., Zalka, C.: Bounds for small-error and zero-error quantum algorithms. In: *Proceedings of the 40th IEEE FOCS*, pp. 358–368 (1999) (arXiv:cs/9904019[cs.CC])
- [BS04] Barnum, H., Saks, M.: A lower bound on the quantum query complexity of read-once functions. *J. Comput. Syst. Sci.* **69**(2), 244–258 (2004). (arXiv:quant-ph/0201007)
- [BSS03] Barnum, H., Saks, M., Szegedy, M.: Quantum query complexity and semi-definite programming. In: *Proceedings of the 18th IEEE, Complexity*, pp. 179–193 (2003)
- [CNW10] Chiang, C.-F., Nagaj, D., Wocjan, P.: Efficient circuits for quantum walks. *Quantum Inf. Comput.* **10**(5–6), 420–434 (2010) (arXiv:0903.3465 [quant-ph])
- [FGG08] Farhi, E., Goldstone, J., Gutmann, S.: A quantum algorithm for the Hamiltonian NAND tree. *Theory Comput.* **4**, 169–190 (2008). (arXiv:quant-ph/0702144)

- [GR02] Grover, L.K., Rudolph, T.: Creating superpositions that correspond to efficiently integrable probability distributions (2002) (arXiv:quant-ph/0208112)
- [Gro96] Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the 28th ACM STOC, pp. 212–219 (1996) (arXiv:quant-ph/9605043)
- [Gro02] Grover, L.K.: Tradeoffs in the quantum search algorithm (2002) (arXiv:quant-ph/0201152)
- [HLŠ05] Høyer, P., Lee, T., Špalek, R.: Tight adversary bounds for composite functions (2005) (arXiv:quant-ph/0509067)
- [HLŠ06] Høyer, P., Lee, T., Špalek, R.: Source codes of semidefinite programs for ADV^\pm . [http://www.ucw.cz/robert/papers/adv/\(2006\)](http://www.ucw.cz/robert/papers/adv/(2006))
- [HLŠ07] Høyer, P., Lee, T., Špalek, R.: Negative weights make adversaries stronger. In: Proceedings of the 39th ACM STOC, pp. 526–535 (2007) (arXiv:quant-ph/0611054)
- [HMW03] Høyer, P., Mosca, M., de Wolf, R.: Quantum search on bounded-error inputs. ICALP 2003. LNCS, vol. 2719, pp. 291–299. Springer, Heidelberg (2003)
- [HNS02] Høyer, P., Neerbek, J., Shi, Y.: Quantum complexities of ordered searching, sorting, and element distinctness. *Algorithmica* **34**(4), 429–448 (2002). (arXiv:quant-ph/0102078. Special issue on Quantum Computation and Cryptography)
- [HNW93] Heiman, Rafi, Newman, Ilan, Wigderson, Avi: On read-once threshold formulae and their randomized decision tree complexity. *Theoret. Comput. Sci.* **107**(1), 63–76 (1993)
- [HW91] Heiman, R., Wigderson, A.: Randomized vs. deterministic decision tree complexity for read-once boolean functions. *Comput. Complex.* **1**(4), 311–329 (1991). (Earlier version in *Structure in Complexity Theory '91*)
- [JKS03] Jayram, T.S., Kumar, R., Sivakumar, D.: Two applications of information complexity. In: Proceedings of the 35th ACM STOC, pp. 673–682 (2003)
- [KSV02] Kitaev, A.Y., Shen, A.H., Vyalı, M.N.: *Classical and Quantum Computation*, vol. 47 of Graduate Studies in Mathematics. American Mathematical Society, Providence (2002)
- [KW93] Karchmer, M., Wigderson, A.: On span programs. In: Proceedings of the 8th IEEE Symposium on Structure in Complexity Theory, pp. 102–111 (1993)
- [LLS06] Laplante, S., Lee, T., Szegedy, M.: The quantum adversary method and classical formula size lower bounds. *Comput. Complex.* **15**, 163–196 (2006). (arXiv:quant-ph/0501057. Earlier version in *Complexity'05*)
- [LM04] Laplante, S., Magniez, F.: Lower bounds for randomized and quantum query complexity using Kolmogorov arguments. In: Proceedings of the 19th IEEE, Complexity, pp. 294–304 (2004) (arXiv:quant-ph/0311189)
- [NC00] Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge (2000)
- [Rei09] Reichardt, B.W.: Span programs and quantum query complexity: the general adversary bound is nearly tight for every boolean function. Extended abstract in: Proceedings of the 50th IEEE FOCS, pp. 544–551 (2009) (arXiv:0904.2759[quant-ph])
- [Rei11] Reichardt, B.W.: Faster quantum algorithm for evaluating game trees. In: Proceedings of the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 546–559 (2011) (arXiv:0907.1623[quant-ph])

- [RŠ08] Reichardt, B.W., Špalek, R.: Span-program-based quantum algorithm for evaluating formulas. In: Proceedings of the 40th ACM STOC, pp. 103–112 (2008) (arXiv:0710.2630[quant-ph])
- [San95] Santha, M.: On the Monte Carlo decision tree complexity of read-once formulae. *Random Struct. Algorithms* **6**(1):75–87 (1995) (Earlier version in Proc. 6th IEEE Structure in Complexity Theory, 1991)
- [Sni85] Snir, M.: Lower bounds on probabilistic linear decision trees. *Theor. Comput. Sci.* **38**, 69–82 (1985)
- [ŠS06] Špalek, R., Szegedy, M.: All quantum adversary methods are equivalent. *Theor. Comput.* **2**(1):1–18 (2006) (arXiv:quant-ph/0409116. Earlier version in ICALP’05)
- [SW86] Saks, M., Wigderson, A.: Probabilistic Boolean decision trees and the complexity of evaluating game trees. In: Proceedings of the 27th IEEE FOCS, pp. 29–38 (1986)
- [Sze04] Szegedy, M.: Quantum speed-up of Markov chain based algorithms. In: Proceedings of the 45th IEEE FOCS, pp. 32–41 (2004)
- [Zha05] Zhang, S.: On the power of Ambainis’s lower bounds. *Theor. Comput. Sci.* **339**(2–3):241–256 (2005) (arXiv:quant-ph/0311060. Earlier version in ICALP’04)