# Faster quantum algorithm for evaluating game trees 

Ben W. Reichardt*


#### Abstract

We give an $O(\sqrt{n} \log n)$-query quantum algorithm for evaluating size- $n$ AND-OR formulas. Its running time is poly-logarithmically greater after efficient preprocessing. Unlike previous approaches, the algorithm is based on a quantum walk on a graph that is not a tree. Instead, the algorithm is based on a hybrid of direct-sum span program composition, which generates tree-like graphs, and a novel tensor-product span program composition method, which generates graphs with vertices corresponding to minimal zero-certificates.

For comparison, by the general adversary bound, the quantum query complexity for evaluating a size- $n$ read-once AND-OR formula is at least $\Omega(\sqrt{n})$, and at most $O(\sqrt{n} \log n / \log \log n)$. However, this algorithm is not necessarily time efficient; the number of elementary quantum gates applied between input queries could be much larger. Ambainis et al. have given a quantum algorithm that uses $\sqrt{n} 2^{O(\sqrt{\log n})}$ queries, with a poly-logarithmically greater running time.


## 1 Introduction

An AND-OR formula is a rooted tree in which the internal nodes correspond to AND or OR gates. The size of the formula is the number of leaves. To a formula $\varphi$ of size $n$ and a numbering of the leaves from 1 to $n$ corresponds a function $\varphi:\{0,1\}^{n} \rightarrow\{0,1\}$. This function is defined on input $x \in\{0,1\}^{n}$ by placing bit $x_{j}$ on the $j$ th leaf, for $j=1,2, \ldots, n$, and evaluating the gates toward the root. Evaluating an AND-OR formula solves the decision version of a MIN-MAX tree, also known as a two-player game tree.

Let $Q(\varphi)$ be the quantum query complexity for evaluating the size- $n$ AND-OR formula $\varphi$. Quantum query complexity is the generalization of classical decision tree complexity to quantum algorithms. Now the general adversary bound of $\varphi$ is $\operatorname{Adv}^{ \pm}(\varphi)=\sqrt{n}[\operatorname{BSO4}$, HLŠ07], and thus $Q(\varphi)=\Omega\left(\operatorname{Adv}^{ \pm}(\varphi)\right)=\Omega(\sqrt{n})$. Since the general adversary bound is nearly tight for any boolean function, in particular $Q(\varphi)=O(\sqrt{n} \log n / \log \log n)$ [Rei09a]. (Interpreted in a different way, this says that the square of the quantum query complexity of evaluating a boolean function is almost a lower bound on the read-once formula size for that function [LLS06].) However, the algorithm from [Rei09a] is not necessarily even time efficient. That is, even though the number of queries to the input is nearly optimal, the number of elementary quantum gates applied between input queries could be much larger.

Ambainis et al. [ $\left.\mathrm{ACR}^{+} 07\right]$ have given a quantum algorithm that evaluates $\varphi$ using $\sqrt{n} 2^{O(\sqrt{\log n})}$ queries, with a running time only poly-logarithmically larger after efficient preprocessing. We reduce the query overhead from $2^{O(\sqrt{\log n})}$ to only $O(\log n)$, with the same preprocessing assumption.

[^0]Theorem 1.1. Let $\varphi$ be an AND-OR formula of size $n$. Then $\varphi$ can be evaluated with error at most $1 / 3$ by a quantum algorithm that uses $O(\sqrt{n} \log n)$ 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 $\sqrt{n}(\log n)^{O(1)}$.

An improvement from $2^{O(\sqrt{\log n})}$ to $O(\log n)$ overhead may not be significant for eventual practical applications. Additionally, the algorithm does not obviously bring us closer to knowing whether the general adversary bound is tight for quantum query complexity, because its overhead is larger than the $O(\log n / \log \log n)$ overhead of the query algorithm from [Rei09a]. (It may be that the analysis used to prove Theorem 1.1 is somewhat loose.)

However, the idea behind the algorithm is still of interest, as it provides a new solution to the problem of evaluating AND-OR formulas with large depth. If $\varphi$ is a formula on $n$ variables, with depth $d$, then the $\left[\mathrm{ACR}^{+} 07\right]$ algorithm, applied directly, evaluates $\varphi$ using $O(\sqrt{n d})$ queries. ${ }^{1}$ Thus for a highly unbalanced formula, with depth $d=\Omega(n)$, the quantum algorithm performs no better asymptotically than the trivial $n$-query classical algorithm. Fortunately, Bshouty et al. have given a "rebalancing" procedure that takes AND-OR formula $\varphi$ as input and outputs an equivalent ANDOR formula $\varphi^{\prime}$ with depth $d^{\prime}=2^{O(\sqrt{\log n})}$ and size $n^{\prime}=n 2^{O(\sqrt{\log n})}$ [BCE91, BB94]. Appealing to this result, $\left[\mathrm{ACR}^{+} 07\right]$ evaluates $\varphi^{\prime}$ using $O\left(\sqrt{n^{\prime} d^{\prime}}\right)=\sqrt{n} 2^{O(\sqrt{\log n})}$ queries.

The algorithm behind Theorem 1.1 gets around the large-depth problem without using formula rebalancing. Instead, the algorithm is based on a novel method for constructing bipartite graphs with certain useful spectral properties. Ambainis et al. run a quantum walk on a graph that matches the formula tree, except with certain edge weights and boundary conditions at the leaves. This tree comes from glueing together elementary graphs for each gate. We term this composition method "direct-sum" composition, because the graph's adjacency matrix acts on a space that is the direct sum of spaces for each individual gate. Direct-sum composition incurs severe overhead on highly unbalanced formulas, making the query complexity at least proportional to the formula depth.

The new algorithm begins with the same elementary graphs, with even the same edge weights. However, it composes them using a kind of "tensor-product" graph composition method. Overall, this results in graphs that have much lower depth, although they are not trees. By carefully combining this method with direct-sum composition, we obtain a graph on which the algorithm runs a quantum walk. The two approaches are summarized in more detail in Section 1.1 below.

The general formula-evaluation problem is an ideal example of a recursively defined problem. The evaluation of a formula is the evaluation of a function, the inputs of which are themselves independent formulas. As argued in a companion paper [Rei09b], quantum computers are particularly well suited for evaluating formulas. Unlike the situation for classical algorithms, for a large class of formulas the optimal quantum algorithm can work following the formula's recursive structure. Since Theorem 1.1 does not require AND-OR formula rebalancing, it extends this quantum recursion paradigm. Besides its conceptual appeal, this may also be important because the effect of rebalancing on the general quantum adversary bound appears to be less tractable for formulas over gate sets beyond AND and OR. Therefore the rebalancing step that aids [ACR $\left.{ }^{+} 07\right]$ might not be useful to solve the large-depth problem on more general formulas. The hybrid graph composition method is another tool that might generalize more easily.

[^1]| Formula $\varphi$ | Randomized, zero-error query complexity $R(\varphi)$ | Quantum bounded-error query complexity $Q(\varphi)$ |
| :---: | :---: | :---: |
| $\mathrm{OR}_{n}$ | $n$ | $\Theta(\sqrt{n}) \quad$ [Gro96, BBBV97] |
| Balanced $\mathrm{AND}_{2}-\mathrm{OR}_{2}$ | $\Theta\left(n^{\alpha}\right) \quad[$ SW86] | $\Theta(\sqrt{n})$ [FGG07, ACR $\left.{ }^{+} 07\right]$ |
| Well-balanced AND-OR | tight recursion [SW86] |  |
| Approx.-balanced AND-OR |  | $\Theta(\sqrt{n}) \quad\left[\mathrm{ACR}^{+} 07\right.$, Rei09b] |
| Arbitrary AND-OR | $\Omega\left(n^{0.51}\right) \quad$ [HW91] | $\Omega(\sqrt{n})$ $[\mathrm{BSO4]}$ <br> $O(\sqrt{n} \log n)$ $($ Thm. 1.1 $)$ |
| Balanced MAJ ${ }_{3}\left(n=3^{d}\right)$ | $\Omega\left((7 / 3)^{d}\right), O\left(2.654^{d}\right)$ [JKS03] | $\Theta\left(2^{d}\right) \quad[\mathrm{RS} 08]$ |
| Balanced over $\mathcal{S}$ |  | $\Theta\left(\operatorname{Adv}^{ \pm}(\varphi)\right) \quad[\mathrm{Rei09a}]$ |
| Almost-balanced over $\mathcal{S}$ |  | $\Theta\left(\operatorname{Adv}^{ \pm}(\varphi)\right) \quad$ [Rei09b] |

Table 1: Comparison of some classical and quantum query complexity results for formula evaluation. Here the exponent $\alpha$ is given by $\alpha=\log _{2}\left(\frac{1+\sqrt{33}}{4}\right) \approx 0.753$, and $\mathcal{S}$ is any fixed, finite gate set. Under certain assumptions, the algorithms' running times are only poly-logarithmically slower.

Our algorithm is developed and analyzed using the framework relating span programs and quantum algorithms from [Rei09a]. The connection to time-efficient quantum algorithms, especially for evaluating unbalanced formulas over arbitrary fixed, finite gate sets, has been developed further in [Rei09b]. Table 1 summarizes some results for the formula-evaluation problem in the classical and quantum models; for a more detailed and inclusive review, see [Rei09b]. Section 1.2 below will go over only the history of quantum algorithms for evaluating AND-OR formulas.

### 1.1 Idea of the algorithm

As an example to illustrate the main idea of our algorithm, consider the AND-OR formula $\varphi(x)=$ $\left(\left[\left(x_{1} \wedge x_{2}\right) \vee x_{3}\right] \wedge x_{4}\right) \vee\left(x_{5} \wedge\left[x_{6} \vee x_{7}\right]\right)$, where $\wedge$ and $\vee$ denote AND and OR, respectively. In Figure 1(a), this formula is represented as a tree.

The $\left[\mathrm{ACR}^{+} 07\right]$ algorithm starts with the graph in Figure 1(b), essentially the same as the formula tree, except with extra edges attached to the root and some leaves. The edges should be weighted, but for this intuitive discussion take every edge's weight to be one.

Consider an input $x \in\{0,1\}^{7}$. Modify the graph by attached a dangling edge to vertex $j$ if $x_{j}=0$, for $j=1,2, \ldots, 7$. Then it is simple to see that the resulting graph has an eigenvalue-zero eigenvector supported on vertex 0 (or rather its adjacency matrix does) if and only if $\varphi(x)=1$. If we added an edge off vertex 0 , then the resulting graph would have an eigenvalue-zero eigenvector supported on the new root if and only if $\varphi(x)=0$.

The $\left[\mathrm{ACR}^{+} 07\right]$ algorithm takes advantage of this property by running (phase estimation on) a quantum walk that starts at the root vertex. The algorithm detects the eigenvalue-zero eigenvector in order to evaluate the formula.

This algorithm does not work well on formulas with large depth. For example, consider the maximally unbalanced formula on $n$ inputs, a skew tree. The corresponding graph is nearly the length- $n$ line graph. It will still have an eigenvalue-zero eigenvector supported on the root if and only if the formula evaluates to zero. However, the algorithm will require $\Omega(n)$ time to detect this eigenvector, because its squared support on the root is only $O(1 / n)$ after normalization, and because


Figure 1: In (a), the AND-OR formula $\varphi(x)=\left(\left[\left(x_{1} \wedge x_{2}\right) \vee x_{3}\right] \wedge x_{4}\right) \vee\left(x_{5} \wedge\left[x_{6} \vee x_{7}\right]\right)$ is represented as a tree, with some gates marked for future reference. Figure (b) shows the graph used by [ACR $\left.{ }^{+} 07\right]$.
the spectral gap around zero will also be $O(1 / n)$. (The spectral gap determines the precision of the phase estimation procedure, and hence its running time. It corresponds to the squared support of eigenvalue-zero eigenvectors by [Rei09a, Theorem 8.7].) Alternatively, one can argue that the algorithm requires $\Omega(n)$ time because it takes that long even to reach the deepest leaf vertices.

Now consider the graph in Figure 2(a). Again modify the graph according to an input $x \in\{0,1\}^{7}$ by attaching dangling edges to those vertices with $x_{j}=0$. Considering a few examples should convince the reader that the resulting graph has an eigenvalue-zero eigenvector supported on vertex 0 if and only if $\varphi(x)=0$. Note, though, that the distance from "output" vertex 0 to any of the "input" vertices 1 to 7 is at most two. The graph is also far from being a tree. Its main feature is that the "constraint" vertices - the vertices aside from $0,1, \ldots, 7$-are in one-to-one correspondence with minimal zero-certificates to $\varphi$. For example, for $x=0101011, \varphi(x)=0$, but flipping any bit of $x$ from 0 to 1 changes the formula evaluation from 0 to 1 . The corresponding constraint vertex is connected to exactly those input vertices $j$ for which $x_{j}=0$.

The graphs in Figures 1(b) and 2(a) represent two extremes of a family of graphs that evaluate $\varphi$ in the same manner. The graph in Figure 1(b) can be seen (essentially) as coming from plugging together along $\varphi$ the individual graphs in Figures 3(a) and 3(b) that evaluate single OR and AND gates in the same manner. We term this "direct-sum" composition. The graph in Figure 2(a) comes from a certain "tensor-product" composition of the graphs from Figures 3(a) and 3(b). (Figure 3 shows several other examples of tensor-product graph composition.)

With the correct choice of edge weights, tensor-product composition leads to graphs for which the squared support of the unit-normalized eigenvalue-zero eigenvector on the root is $\Omega(1 / \sqrt{n})$ for a formula of size $n$. This implies by [Rei09a, Theorem 9.1] a quantum algorithm that uses $O(\sqrt{n} \log n / \log \log n)$ queries to evaluate the formula. There is no issue with deep formulas. However, the algorithm will not be time efficient. Essentially, the problem is that the number of vertices can be exponentially large in $n$, as can be their degrees, which makes it difficult to implement a quantum walk efficiently.

Theorem 1.1 gets around this problem by using a combination of the two methods of composing graphs. For example, the graph in Figure 2(b) also evaluates the formula $\varphi$ in the same manner. One can think of the vertex $c$ as evaluating the subformula $x_{5} \wedge\left(x_{6} \vee x_{7}\right)$. The subgraph it cuts off


Figure 2: For the formula from Figure 1(a), we display the reduced tensor-product composed span program in (a), and the hybrid-composed span program in (b) for the case when the edge to the subformula on inputs $x_{5}, x_{6}, x_{7}$ is checkpointed (see Section 3.3). In each case, the output vertex corresponding to the target vector is labeled 0 and the input bits are labeled from 1 to 7. The other vertices in (a) are labeled according to the corresponding maximal false inputs (see Definition 3.5 and Lemma 3.6). In (b), the checkpoint vertex is labeled c-this vertex corresponds to the free input vector added in direct-sum composition (see Definition 2.4) -and the other vertices have been labeled according to maximal false inputs with the checkpointed subformula grouped together. In (c), the edge weights for the graph in (b) are specified using the biadjacency matrix. The parameters $\alpha, \delta, \ldots$ are marked with ' or " to indicate to which gate in the formula of Figure 1(a) they correspond.
has been composed in a direct-sum manner with a tensor-product-composed graph for the formula $\left(\left[\left(x_{1} \wedge x_{2}\right) \vee x_{3}\right] \wedge x_{4}\right) \vee x_{c}$. By combining the two composition methods, we can manage the tradeoffs, controlling the maximum degree and norm of the graph, while also avoiding the formula depth problem.

Although our algorithm can be presented and analyzed entirely in terms of graphs, we will present it in terms of span programs. Span programs are part of a framework for designing and analyzing quantum algorithms [Rei09a], for which Section 2 gives some necessary background. Eigenvalue-zero eigenvectors correspond to span program witnesses and the squared support on the root corresponds to a complexity measure known as the full witness size [Rei09b]. The two graph composition techniques described above correspond to different ways of composing general span programs.

### 1.2 Review of quantum algorithms for evaluating AND-OR formulas

Research on the formula-evaluation problem in the quantum model began with the simple $n$-bit OR function, $\mathrm{OR}_{n}$. Grover gave a quantum algorithm for evaluating $\mathrm{OR}_{n}$ with bounded one-sided error using $O(\sqrt{n})$ oracle queries and $O(\sqrt{n} \log \log n)$ time [Gro96, Gro02].

Grover's algorithm, together with repetition for reducing errors, can be applied recursively to speed up the evaluation of more general AND-OR formulas. For example, the size- $n$ ANDOR formula $\mathrm{AND}_{\sqrt{n}} \circ\left(\mathrm{OR}_{\sqrt{n}}, \ldots, \mathrm{OR}_{\sqrt{n}}\right)$ can be evaluated in $O(\sqrt{n} \log n)$ queries. Here the extra logarithmic factor comes from using repetition to reduce the error probability of the inner OR $_{\sqrt{n}}$ evaluation procedure from a constant to be polynomially small. Call a formula layered if the gates at the same depth are the same. Buhrman, Cleve and Wigderson show that the above argument can be applied to evaluate a layered, depth- $d$ AND-OR formula on $n$ inputs using $O\left(\sqrt{n} \log ^{d-1} n\right)$ queries [BCW98, Theorem 1.15].

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

$$
\begin{equation*}
\tilde{O}_{x}:|\varphi\rangle \otimes|j\rangle \otimes|b\rangle \otimes|0\rangle \mapsto|\varphi\rangle \otimes|j\rangle \otimes\left(\left|b \oplus x_{j}\right\rangle \otimes\left|\psi_{x, j, x_{j}}\right\rangle+\left|b \oplus \bar{x}_{j}\right\rangle \otimes\left|\psi_{x, j, \bar{x}_{j}}\right\rangle\right), \tag{1.1}
\end{equation*}
$$

where $\left|\psi_{x, j, x_{j}}\right\rangle,\left|\psi_{x, j, \bar{x}_{j}}\right\rangle$ are pure states with $\|\left|\psi_{x, j, x_{j}}\right\rangle \|^{2} \geq 2 / 3$. Such an oracle can be implemented when the function $j \stackrel{\rightharpoonup}{~_{j}}$ is computed by a bounded-error, randomized subroutine [NC00]. Høyer et al. allow access to $\tilde{O}_{x}$ and $\tilde{O}_{x}^{-1}$, both at unit cost, and show that $\mathrm{OR}_{n}$ 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 $O\left(\sqrt{n} c^{d-1}\right)$ queries, for some constant $c>1000$. This gives an $O(\sqrt{n})$-query quantum algorithm for the case that the depth $d$ is constant, but is not sufficient to cover, e.g., the complete, binary AND-OR formula, for which $d=\log _{2} n$.

A breakthrough for the formula-evaluation problem came in 2007, when Farhi, Goldstone and Gutmann presented a quantum algorithm for evaluating complete, binary AND-OR formulas [FGG07]. Their algorithm is not based on iterating Grover's algorithm in any way, but instead runs a quantum walk - analogous to a classical random walk - on a graph derived from the ANDOR formula graph as in Figure 1. The algorithm runs in time $O(\sqrt{n})$ in a certain continuous-time query model.

Ambainis et al. discretized the [FGG07] 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 $\left[\mathrm{ACR}^{+} 07\right]$. Applying this correspondence to
quantum walks on certain weighted graphs, they gave an $O(\sqrt{n})$-query quantum algorithm for evaluating "approximately balanced" formulas, extended in [Rei09b]. Using the formula rebalancing procedure of [BCE91, BB94], the [ACR $\left.{ }^{+} 07\right]$ algorithm uses $\sqrt{n} 2^{O(\sqrt{\log n})}$ queries in general. This is nearly optimal, since the adversary bound gives an $\Omega(\sqrt{n})$ lower bound [BS04].

This author has given an $O(\sqrt{n} \log n / \log \log n)$-query quantum algorithm for evaluating arbitrary size- $n$ AND-OR formulas [Rei09a]. In fact, the result is more general, stating that the general adversary bound is nearly tight for every boolean function. However, unlike the earlier AND-OR formula-evaluation algorithms, the algorithm is not necessarily time efficient.

## 2 Span programs

In this section, we will briefly recall some of the definitions and results on span programs from [Rei09a, Rei09b]. This section is essentially an abbreviated version of [Rei09b, Sec. 2].

For a natural number $n$, let $[n]=\{1,2, \ldots, n\}$. 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 its operator norm. Let $B=\{0,1\}$. For a string $x \in B^{n}$, let $\bar{x}$ denote its bitwise complement.

### 2.1 Span program full witness size

The full witness size is a span program complexity measure that is important for developing quantum algorithms that are time efficient as well as query efficient.

Definition 2.1 (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 $I_{\text {free }}$ and $I_{j, b}$ for $j \in[n], b \in B$, and "input vectors" $\left|v_{i}\right\rangle \in V$ for $i \in I_{\text {free }} \cup \bigcup_{j \in[n], b \in B} I_{j, b}$.

To $P$ corresponds a function $f_{P}: B^{n} \rightarrow B$, defined on $x \in B^{n}$ by

$$
f_{P}(x)= \begin{cases}1 & \text { if }|t\rangle \in \operatorname{Span}\left(\left\{\left|v_{i}\right\rangle: i \in I_{\text {free }} \cup \bigcup_{j \in[n]} I_{j, x_{j}}\right\}\right)  \tag{2.1}\\ 0 & \text { otherwise }\end{cases}
$$

Some additional notation is convenient. Fix a span program $P$. Let $I=I_{\text {free }} \cup \bigcup_{j \in[n], b \in B} I_{j, b}$. Let $A \in \mathcal{L}\left(\mathbf{C}^{I}, V\right)$ be given by $A=\sum_{i \in I}\left|v_{i}\right\rangle\langle i|$. For $x \in B^{n}$, let $I(x)=I_{\text {free }} \cup \bigcup_{j \in[n]} I_{j, x_{j}}$ and $\Pi(x)=\sum_{i \in I(x)}|i\rangle\langle i| \in \mathcal{L}\left(\mathbf{C}^{I}\right)$. Then $f_{P}(x)=1$ if $|t\rangle \in \operatorname{Range}(A \Pi(x))$. 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 $\left|w^{\prime}\right\rangle \in V$ is said to be a witness for $f_{P}(x)=0$ if $\left\langle t \mid w^{\prime}\right\rangle=1$ and $\Pi(x) A^{\dagger}\left|w^{\prime}\right\rangle=0$.
Definition 2.2 (Witness size). Consider a span program $P$, and a vector $s \in[0, \infty)^{n}$ of nonnegative "costs." Let $S=\sum_{j \in[n], b \in B, i \in I_{j, b}} \sqrt{s_{j}}|i\rangle\langle i| \in \mathcal{L}\left(\mathbf{C}^{I}\right)$. For each input $x \in B^{n}$, define the witness size of $P$ on $x$ with costs $s$, wsize $_{s}(P, x)$, as follows:

$$
\operatorname{wsize}_{s}(P, x)= \begin{cases}\min _{|w\rangle: A \Pi(x)|w\rangle=|t\rangle} \| S|w\rangle \|^{2} & \text { if } f_{P}(x)=1  \tag{2.2}\\ \min _{\left|w^{\prime}\right\rangle:\left\langle t \mid w^{\prime}\right\rangle=1} \| S A^{\dagger}\left|w^{\prime}\right\rangle \|^{2} & \text { if } f_{P}(x)=0 \\ \Pi(x) A^{\dagger}\left|w^{\prime}\right\rangle=0\end{cases}
$$

The witness size of $P$ with costs $s$ is

$$
\begin{equation*}
\operatorname{wsize}_{s}(P)=\max _{x \in B^{n}} \operatorname{wsize}_{s}(P, x) \tag{2.3}
\end{equation*}
$$

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

$$
\begin{align*}
\operatorname{fwsize}_{s}(P, x) & = \begin{cases}\min _{|w\rangle: A \Pi(x)|w\rangle=|t\rangle}\left(1+\| S^{f}|w\rangle \|^{2}\right) & \text { if } f_{P}(x)=1 \\
\left.\min _{\left|w^{\prime}\right\rangle:\left\langle t \mid w^{\prime}\right\rangle=1}\left(\|| | w^{\prime}\right\rangle\left\|^{2}+\right\| S A^{\dagger}\left|w^{\prime}\right\rangle \|^{2}\right) & \text { if } f_{P}(x)=0 \\
\Pi(x) A^{\dagger}\left|w^{\prime}\right\rangle=0\end{cases}  \tag{2.4}\\
\text { fwsize }_{s}(P) & =\max _{x \in B^{n}} \operatorname{fwsize}_{s}(P, x) . \tag{2.5}
\end{align*}
$$

When the subscript $s$ is omitted, the costs are taken to be uniform, $s=\overrightarrow{1}=(1,1, \ldots, 1)$, e.g., fwsize $(P)=\operatorname{fwsize}_{\overrightarrow{\mathrm{r}}}(P)$. The witness size is defined in [RŠ08]. The full witness size is defined in [Rei09a, Sec. 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 $j$ [Rei09a, Def. 4.9].

Theorem 2.3 ([Rei09a, Theorem 9.3], [Rei09b, Theorem 2.3]). Let $P$ be a span program. Then $f_{P}$ can be evaluated using

$$
\begin{equation*}
T=O\left(\operatorname{fwsize}(P)\left\|\operatorname{abs}\left(A_{G_{P}}\right)\right\|\right) \tag{2.6}
\end{equation*}
$$

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.

### 2.2 Direct-sum span program composition

Let $f: B^{n} \rightarrow B$ be a boolean function. Let $S \subseteq[n]$. For $j \in[n]$, let $m_{j}$ be a natural number, with $m_{j}=1$ for $j \notin S$. For $j \in S$, let $f_{j}: B^{m_{j}} \rightarrow B$. Define $y: B^{m_{1}} \times \cdots \times B^{m_{n}} \rightarrow B^{n}$ by

$$
y(x)_{j}= \begin{cases}f_{j}\left(x_{j}\right) & \text { if } j \in S  \tag{2.7}\\ x_{j} & \text { if } j \notin S\end{cases}
$$

Define $g: B^{m_{1}} \times \cdots \times B^{m_{n}} \rightarrow B$ by $g(x)=f(y(x))$. Given span programs for the individual functions $f$ and $f_{j}$ for $j \in S$, we will construct a span program for $g$.

Let $P$ be a span program computing $f_{P}=f$. Let $P$ have inner product space $V$, target vector $|t\rangle$ and input vectors $\left|v_{i}\right\rangle$ indexed by $I_{\text {free }}$ and $I_{j c}$ for $j \in[n]$ and $c \in B$.

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 $P^{j 0}$ and $P^{j 1}$ be span programs computing $f_{P^{j 1}}=f_{j}: B^{m_{j}} \rightarrow B$ and $f_{P^{j 0}}=\neg f_{j}$, with $r_{j}=\operatorname{wsize}_{s_{j}}\left(P^{j 0}\right)=\operatorname{wsize}_{s_{j}}\left(P^{j 1}\right)$. For $c \in B$, let $P^{j c}$ have inner product space $V^{j c}$ with target vector $\left|t^{j c}\right\rangle$ and input vectors indexed by $I_{\text {free }}^{j c}$ and $I_{k b}^{j c}$ for $k \in\left[m_{j}\right], b \in B$. For $j \notin S$, let $r_{j}=s_{j}$.

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

Definition 2.4 ([Rei09a, Def. 4.5]). The direct-sum-composed span program $Q^{\oplus}$ is defined by:

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

$$
\left|v_{i}^{\oplus}\right\rangle= \begin{cases}\left|v_{i}\right\rangle_{V} & \text { if } i \in I_{\text {free }}  \tag{2.8}\\ \left|v_{i}\right\rangle_{V}+|i\rangle \otimes\left|t^{j c}\right\rangle & \text { if } i \in I_{j c} \text { and } j \in S \\ \left|i^{\prime}\right\rangle \otimes\left|v_{i^{\prime \prime}}\right\rangle & \text { if } i=\left(i^{\prime}, i^{\prime \prime}\right) \in I_{j c} \times I_{\text {free }}^{j c}\end{cases}
$$

- The other input vectors are indexed by $I_{(j k) b}^{\oplus}$ for $j \in[n], k \in\left[m_{j}\right], b \in B$. For $j \notin S$, $I_{(j 1) b}^{\oplus}=I_{j b}$, with $\left|v_{i}^{\oplus}\right\rangle=\left|v_{i}\right\rangle_{V}$ for $i \in I_{(j 1) b}^{\oplus}$. For $j \in S$, let $I_{(j k) b}^{\oplus}=\bigcup_{c \in B}\left(I_{j c} \times I_{k b}^{j c}\right)$. For $i \in I_{j c}$ and $i^{\prime} \in I_{k b}^{j c}$, let

$$
\begin{equation*}
\left|v_{i i^{\prime}}^{\oplus}\right\rangle=|i\rangle \otimes\left|v_{i^{\prime}}\right\rangle . \tag{2.9}
\end{equation*}
$$

By [Rei09a, Theorem 4.3], $f_{Q^{\oplus}}=g$ and $\operatorname{wsize}_{s}\left(Q^{\oplus}\right) \leq \operatorname{wsize}_{r}(P)$. Moreover, we can bound how quickly the full witness size can grow relative to the witness size:

Lemma 2.5 ([Rei09b, Lemma 2.5]). Under the above conditions, for each input $x \in B^{m_{1}} \times \cdots \times$ $B^{m_{n}}$, with $y=y(x)$,

- If $g(x)=1$, let $|w\rangle$ be a witness to $f_{P}(y)=1$ with $\sum_{j \in[n], i \in I_{j y_{j}}} r_{j}\left|w_{i}\right|^{2}=\operatorname{wsize}_{r}(P, y)$. Then

$$
\begin{align*}
& \frac{\operatorname{fwsize}_{s}\left(Q^{\oplus}, x\right)}{\operatorname{wsize}_{r}(P, y)} \leq \sigma(y,|w\rangle)+\frac{1+\sum_{i \in I_{\text {free }}\left|w_{i}\right|^{2}}}{\operatorname{wsize}_{r}(P, y)} \\
& \text { where } \sigma(y,|w\rangle)=\max _{\exists i \in I_{j y_{j}} \in S:}^{\substack{\text { with }\langle i \mid w\rangle \neq 0}} \frac{\text { fwsize }_{s_{j}}\left(P^{j y_{j}}\right)}{\operatorname{wsize}_{s_{j}}\left(P^{j y_{j}}\right)} . \tag{2.10}
\end{align*}
$$

- If $g(x)=0$, let $\left|w^{\prime}\right\rangle$ be a witness to $f_{P}(y)=0$ such that $\sum_{j \in[n], i \in I_{j \bar{y}_{j}}} r_{j}\left|\left\langle w^{\prime} \mid v_{i}\right\rangle\right|^{2}=\operatorname{wsize}_{r}(P, y)$. Then

$$
\begin{align*}
& \frac{\operatorname{fwsize}_{s}\left(Q^{\oplus}, x\right)}{\operatorname{wsize}_{r}(P, y)} \leq \sigma\left(\bar{y},\left|w^{\prime}\right\rangle\right)+\frac{\|\left|w^{\prime}\right\rangle \|^{2}}{\operatorname{wsize}_{r}(P, y)} \\
& \text { where } \sigma\left(\bar{y},\left|w^{\prime}\right\rangle\right)=\max _{\substack{ \\
j \in S: \\
\exists i \in I_{j \bar{y}_{j}} \\
\text { with }\left\langle v_{i} \mid w^{\prime}\right\rangle \neq 0}} \frac{\operatorname{fwsize}_{s_{j}}\left(P^{j \bar{y}_{j}}\right)}{\operatorname{wsize}_{s_{j}}\left(P^{j \bar{y}_{j}}\right)} \tag{2.11}
\end{align*}
$$

If $S=\emptyset$, then $\sigma(y,|w\rangle)$ and $\sigma\left(\bar{y},\left|w^{\prime}\right\rangle\right)$ should each be taken to be 1 in the above equations.
Lemma 2.6 ([Rei09b, Lemma 3.4]). If $P_{\varphi}$ is the direct-sum composition along a formula $\varphi$ of span programs $P_{v}$ and $P_{v}^{\dagger}$, then

$$
\begin{equation*}
\left\|\operatorname{abs}\left(A_{G_{P}}\right)\right\| \leq 2 \max _{v \in \varphi} \max \left\{\left\|\operatorname{abs}\left(A_{G_{P_{v}}}\right)\right\|,\left\|\operatorname{abs}\left(A_{G_{P_{v}^{\dagger}}}\right)\right\|\right\} . \tag{2.12}
\end{equation*}
$$

If the span programs $P_{v}$ are monotone, then $\left\|\operatorname{abs}\left(A_{G_{P}}\right)\right\| \leq 2 \max _{v}\left\|\operatorname{abs}\left(A_{G_{P_{v}}}\right)\right\|$.

## 3 Evaluation of arbitrary AND-OR formulas

Let $\varphi$ be an AND-OR formula. In this section, we will prove Theorem 1.1 by applying Theorem 2.3 to a certain composed span program $P_{\varphi}$. The construction of $P_{\varphi}$ has three steps that we will explain in detail below.

1. Eliminate any AND or OR gates with fan-in one, and expand out AND and OR gates with higher fan-ins into gates of fan-in exactly two.
2. Mark a certain subset of the edges of the formula. We call marked edges "checkpoints."
3. Starting with the span programs for $\mathrm{AND}_{2}$ and $\mathrm{OR}_{2}$, compose them for the subformulas cut off by checkpointed edges using a version of tensor-product composition. Compose the resulting span programs across checkpoints using direct-sum composition to yield $P_{\varphi}$.
Direct-sum span program composition keeps the norm of the corresponding graph's adjacency matrix under control, as well as the maximum degree of a vertex in the graph. However, it makes the full witness size grow much larger than the witness size, especially for highly unbalanced formulas. Reduced tensor-product composition generates strict span programs, for which the full witness size stays close to the witness size. However, it allows the norm and the maximum degree of the corresponding graph to grow exponentially quickly. By using both techniques in careful combination, we are able to manage this tradeoff so that Theorem 2.3 can be applied.

Section 3.1 presents the span programs we use for fan-in-two AND and OR gates.
In Section 3.2, we study reduced tensor-product composition for the span programs for AND and OR gates. Reduced tensor-product composition is a version of tensor-product composition that parsimoniously uses fewer dimensions when possible. For AND-OR formulas, it has the advantage that the output span program's inner product space bears a close connection to false inputs of the formula $\varphi$, similarly to canonical span programs. In order to motivate the checkpointing idea, we explain the problems of a scheme based only on reduced tensor-product composition.

Section 3.3 then presents our full construction of $P_{\varphi}$, based on a combination of direct-sum and reduced tensor-product composition.

Section 3.4 contains the analysis of the graphs $G_{P_{\varphi}}(x)$ required to apply Theorem 2.3.

### 3.1 Span programs for $\mathrm{AND}_{2}$ and $\mathrm{OR}_{2}$

We will use the following strict, monotone span programs for fan-in-two AND and OR gates:
Definition 3.1 ([Rei09b, Def. 4.1]). For $s_{1}, s_{2}>0$, define span programs $P_{\text {AND }}\left(s_{1}, s_{2}\right)$ and $P_{\mathrm{OR}}\left(s_{1}, s_{2}\right)$ computing $\mathrm{AND}_{2}$ and $\mathrm{OR}_{2}, B^{2} \rightarrow B$, respectively, by

$$
\begin{array}{rlll}
P_{\mathrm{AND}}\left(s_{1}, s_{2}\right): & |t\rangle=\binom{\alpha_{1}}{\alpha_{2}}, & & \left|v_{1}\right\rangle=\binom{\beta_{1}}{0},
\end{array} r\left|v_{2}\right\rangle=\binom{0}{\beta_{2}}
$$

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

$$
\begin{align*}
\alpha_{j} & =\left(s_{j} / s_{p}\right)^{1 / 4} & \beta_{j} & =1  \tag{3.3}\\
\delta & =1 & \epsilon_{j} & =\left(s_{j} / s_{p}\right)^{1 / 4}
\end{align*}
$$

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 $\alpha, \epsilon \in\left(1,2^{1 / 4}\right]$. They are largest when $s_{1}=s_{2}$.
Claim 3.2 ([Rei09b, Claim 4.2]). The span programs $P_{\mathrm{AND}}\left(s_{1}, s_{2}\right)$ and $P_{\mathrm{OR}}\left(s_{1}, s_{2}\right)$ satisfy:

$$
\begin{align*}
& \operatorname{wsize}_{\left(\sqrt{s_{1}}, \sqrt{s_{2}}\right)}\left(P_{\mathrm{AND}}, x\right)= \begin{cases}\sqrt{s_{p}} & \text { if } x \in\{11,10,01\} \\
\frac{\sqrt{s_{p}}}{2} & \text { if } x=00\end{cases} \\
& \operatorname{wsize}_{\left(\sqrt{s_{1}}, \sqrt{s_{2}}\right)}\left(P_{\mathrm{OR}}, x\right)= \begin{cases}\sqrt{s_{p}} & \text { if } x \in\{00,10,01\} \\
\frac{\sqrt{s_{p}}}{2} & \text { if } x=11\end{cases} \tag{3.5}
\end{align*}
$$

### 3.2 Reduced tensor-product span program composition for AND-OR formulas

Reduced tensor-product composition of span programs is introduced in [Rei09a, Def. 4.6]. We repeat the definition here, specializing to the case of monotone, strict span programs acting on disjoint inputs. Also, for simplicity we consider the case of composing on one span program at a time. After stating the definition, we specialize further to AND-OR formulas, and characterize the reduced tensor-product span program composition of the AND and OR span programs from Definition 3.1.

Consider monotone functions $f: B \times B^{n} \rightarrow B$ and $f^{\prime}: B^{m} \rightarrow B$. Let $g: B^{m} \times B^{n} \rightarrow B$ be defined by

$$
\begin{equation*}
g(x, y)=f\left(f^{\prime}(x), y\right) \tag{3.6}
\end{equation*}
$$

for $x \in B^{m}, y \in B^{n}$. Let $P$ be a span program computing $f_{P}=f$ and let $P^{\prime}$ be a span program computing $f_{P^{\prime}}=f^{\prime}$. Assume that $P$ and $P^{\prime}$ are both monotone, strict span programs.

Let span program $P$ be in inner product space $V=\mathbf{C}^{[d]}$, with target vector $|t\rangle$ and input vectors $\left\{\left|v_{i}\right\rangle\right\}$ indexed by $I_{j 1}$ for $j \in[n]$. Let $P^{\prime}$ be in the inner product space $V^{\prime}$ with target vector $\left|t^{\prime}\right\rangle$ and input vectors $\left\{\left|v_{i^{\prime}}^{\prime}\right\rangle\right\}$ indexed by $I_{k 1}^{\prime}$ for $k \in[m]$. Since $P$ and $P^{\prime}$ are both monotone, $I_{j 0}=I_{k 0}^{\prime}=\emptyset$.

Definition 3.3 ([Rei09a]). The tensor-product-composed span program, reduced with respect to the basis $\{|l\rangle: l \in[d]\}$ for $V$, is $Q^{r \otimes}$, defined by:

- Let $Z=\left\{l \in[d]: \forall i \in I_{11},\left\langle l \mid v_{i}\right\rangle=0\right\}$. For $l \in[d]$, define $V_{l}$ and $\left|\pi_{l}\right\rangle \in V_{l}$ by

$$
\begin{align*}
V_{l} & = \begin{cases}V^{\prime} & \text { if } l \notin Z, \text { i.e., }\left\langle l \mid v_{i}\right\rangle \neq 0 \text { for some } i \in I_{11} \\
\mathbf{C} & \text { if } l \in Z\end{cases}  \tag{3.7}\\
\left|\pi_{l}\right\rangle & = \begin{cases}\left|t^{\prime}\right\rangle & \text { if } l \notin Z \\
\|\left|t^{\prime}\right\rangle \| & \text { if } l \in Z\end{cases}
\end{align*}
$$

- The inner product space of $Q^{r \otimes}$ is $V^{r \otimes}=\bigoplus_{l \in[d]} V_{l}$. Any vector $|v\rangle \in V^{r \otimes}$ can be uniquely expressed as $\sum_{l \in[d]}\left|v_{l}\right\rangle_{V_{l}}$, where $\left|v_{l}\right\rangle \in V_{l}$.
- The target vector is

$$
\begin{equation*}
\left|t^{r \otimes}\right\rangle=\sum_{l \in[d]}\langle l \mid t\rangle\left|\pi_{l}\right\rangle_{V_{l}} \tag{3.8}
\end{equation*}
$$

- $Q^{r \otimes}$ is strict and monotone, thus $I_{\text {free }}^{r \otimes}=I_{k 0}^{r \otimes}=\emptyset$ for all $k \in[m+n]$.
- The input vectors $\left\{\left|v_{\iota}^{r \otimes}\right\rangle\right\}$ are indexed by

$$
I_{k 1}^{r \otimes}= \begin{cases}I_{11} \times I_{k 1}^{\prime} & \text { if } k \leq m  \tag{3.9}\\ I_{j 1} & \text { if } k>m, \text { where } j=k-m+1\end{cases}
$$

and given by

$$
\left|v_{\iota}^{r \otimes}\right\rangle= \begin{cases}\sum_{l \in[d]}\left\langle l \mid v_{i}\right\rangle\left|v_{i^{\prime}}^{\prime}\right\rangle_{V_{l}} & \text { if } \iota=\left(i, i^{\prime}\right) \in \bigcup_{k \leq m} I_{k 1}^{r \otimes}  \tag{3.10}\\ \sum_{l \in[d]}\left\langle l \mid v_{\iota}\right\rangle\left|\pi_{l}\right\rangle_{V_{l}} & \text { if } \iota \in \bigcup_{k>m} I_{k 1}^{r \otimes}\end{cases}
$$

The intuition behind this construction is to compose the span programs in a tensor-product manner somewhat similar to Definition 2.4. From [Rei09a, Def. 4.4], this would give a span program with target vector $|t\rangle \otimes\left|t^{\prime}\right\rangle \in V \otimes V^{\prime}$ and input vectors either $\left|v_{i}\right\rangle \otimes\left|v_{i^{\prime}}^{\prime}\right\rangle$ for $i \in I_{11}$ or $\left|v_{i}\right\rangle \otimes\left|t^{\prime}\right\rangle$ otherwise. However, if all the $I_{11}$ input vectors are zero in a coordinate $l \in[d]$, then the first type of input vectors are all zero on $|l\rangle \otimes V^{\prime}$. The second type of input vectors are all proportional to the same state $|l\rangle \otimes\left|t^{\prime}\right\rangle$ on that coordinate, so we might as well just keep $|l\rangle \|\left|t^{\prime}\right\rangle \|$ as in the above definition. The advantage over tensor-product composition is that the graph $G_{Q^{r \otimes}}$ associated to the span program $Q^{r \otimes}$ potentially has fewer vertices, with lower degrees.

Here we have composed the span program $P_{1}$ into the first input of $P$. Reduced tensor-product composition into the other inputs is defined symmetrically.

By [Rei09a, Prop. 4.7], for arbitrary costs $r \in[0, \infty)^{m}$ and $s \in[0, \infty)^{n}$,

$$
\begin{equation*}
\operatorname{wsize}_{(r, s)}\left(Q^{r \otimes}\right) \leq \operatorname{wsize}_{\left(\operatorname{wsize}_{r}\left(P^{\prime}\right), s\right)}(P) \tag{3.11}
\end{equation*}
$$

Now let us study reduced tensor-product composition for AND-OR formulas. To start with, it will be helpful to give two examples of Definition 3.3, for the cases $P^{\prime}=P_{\mathrm{OR}}$ and $P^{\prime}=P_{\text {AND }}$. Recall from [Rei09a, Def. 8.1] that the biadjacency matrix for the bipartite graph $G_{P}\left(1^{1+n}\right)$ is given by $B_{G_{P}\left(1^{1+n}\right)}=(|t\rangle \quad A)$, where $A$ is the matrix whose columns are the input vectors of $P$, as defined in Section 2.1. Assume that $I_{11}=\{1\}$ is a singleton set, with $\left|v_{1}\right\rangle$ the first column of A. Rearrange the rows so that the nonzero entries of $\left|v_{1}\right\rangle$ come first (the set $Z$ from Definition 3.3 comes last), writing $\left|v_{1}\right\rangle=(|\gamma\rangle, 0)$, where $|\gamma\rangle$ is nonzero in every entry. Writing $|t\rangle=\left(\left|t_{1}\right\rangle,\left|t_{2}\right\rangle\right)$, $B_{G_{P}\left(1^{1+n}\right)}$ decomposes as

$$
B_{G_{P}\left(1^{1+n}\right)}=\left(\begin{array}{ccc}
\left|t_{1}\right\rangle & |\gamma\rangle & C_{1}  \tag{3.12}\\
\left|t_{2}\right\rangle & 0 & C_{2}
\end{array}\right) .
$$

- First consider the case that $P^{\prime}=P_{\text {AND }}\left(s_{1}, s_{2}\right)$ from Definition 3.1. Let $Q^{r \otimes}$ be the composed span program from Definition 3.3. Then

$$
B_{G_{Q^{r \otimes}}\left(1^{2+n}\right)}=\left(\begin{array}{cccc}
\alpha_{1}\left|t_{1}\right\rangle & \beta_{1}|\gamma\rangle & 0 & \alpha_{1} C_{1}  \tag{3.13}\\
\alpha_{2}\left|t_{1}\right\rangle & 0 & \beta_{2}|\gamma\rangle & \alpha_{2} C_{1} \\
\alpha\left|t_{2}\right\rangle & 0 & 0 & \alpha C_{2}
\end{array}\right)
$$

For comparison, the tensor-product-composed span program $Q^{\otimes}$ from [Rei09a, Def. 4.4] would have the biadjacency matrix

$$
B_{G_{Q} r \otimes}\left(1^{2+n}\right)=\left(\begin{array}{cccc}
\alpha_{1}\left|t_{1}\right\rangle & \beta_{1}|\gamma\rangle & 0 & \alpha_{1} C_{1}  \tag{3.14}\\
\alpha_{2}\left|t_{1}\right\rangle & 0 & \beta_{2}|\gamma\rangle & \alpha_{2} C_{1} \\
\alpha_{1}\left|t_{2}\right\rangle & 0 & 0 & \alpha_{1} C_{2} \\
\alpha_{2}\left|t_{2}\right\rangle & 0 & 0 & \alpha_{2} C_{2}
\end{array}\right) .
$$

- Next consider the case that $P^{\prime}=P_{\mathrm{OR}}\left(s_{1}, s_{2}\right)$ from Definition 3.1. Let $Q^{r \otimes}$ be the composed span program from Definition 3.3. Then

$$
B_{G_{Q^{r} \otimes}\left(1^{2+n}\right)}=\left(\begin{array}{cccc}
\delta\left|t_{1}\right\rangle & \epsilon_{1}|\gamma\rangle & \epsilon_{2}|\gamma\rangle & \delta C_{1}  \tag{3.15}\\
\delta\left|t_{2}\right\rangle & 0 & 0 & \delta C_{2}
\end{array}\right) .
$$

From Eqs. (3.13) and (3.15), we can derive a bound on the growth of the norm of the entrywise absolute value of the biadjacency matrix for a span program, under reduced tensor-product composition with either $P_{\mathrm{AND}}$ or $P_{\mathrm{OR}}$ :

Lemma 3.4. Let $P$ be a strict, monotone span program on $1+n$ input bits, with $\left|I_{1,1}\right|=1$. For $s_{1}, s_{2}>0$, let $P^{\prime}$ be either $P_{\mathrm{AND}}\left(s_{1}, s_{2}\right)$ or $P_{\mathrm{OR}}\left(s_{1}, s_{2}\right)$, from Definition 3.1. Let $Q^{r \otimes}$ be the reduced tensor-product composition of $P^{\prime}$ into the first input of $P$, as in Definition 3.3. Then

$$
\begin{equation*}
\left\|\operatorname{abs}\left(B_{G_{Q^{r} \otimes}\left(1^{2+n}\right)}\right)\right\|^{2} \leq \frac{\sqrt{s_{1}}+\sqrt{s_{2}}}{\sqrt{s_{1}+s_{2}}}\left\|\operatorname{abs}\left(B_{G_{P}\left(1^{1+n}\right)}\right)\right\|^{2} . \tag{3.16}
\end{equation*}
$$

Proof. Indeed, first consider the case that $P^{\prime}=P_{\mathrm{OR}}\left(s_{1}, s_{2}\right)$. By rearranging and regrouping the columns of the biadjacency matrices, from Eq. (3.12), abs $\left(B_{G_{P}\left(1^{1+n}\right)}\right)$ can be rewritten as

$$
\left(\begin{array}{ll}
A & B \tag{3.17}
\end{array}\right)
$$

for some entry-wise nonnegative matrices $A$ and $B$, such that from Eq. (3.15), abs $\left(B_{G_{Q^{r \otimes}}\left(1^{2+n}\right)}\right)$ can be rewritten as

$$
\left(\begin{array}{lll}
\delta A & \epsilon_{1} B & \epsilon_{2} B \tag{3.18}
\end{array}\right) .
$$

Now if $\delta$ were $\sqrt{\epsilon_{1}^{2}+\epsilon_{2}^{2}}$, then the norm of this latter matrix would be exactly $\sqrt{\epsilon_{1}^{2}+\epsilon_{2}^{2}}$ times the norm of the former matrix. Since in fact $\delta=1<\sqrt{\epsilon_{1}^{2}+\epsilon_{2}^{2}}$, we have the desired inequality (3.16).

Next consider the case that $P^{\prime}=P_{\mathrm{OR}}\left(s_{1}, s_{2}\right)$. Rearrange the columns in Eqs. (3.12) and (3.13) to rewrite $\operatorname{abs}\left(B_{G_{P}\left(1^{1+n}\right)}\right)$ and $\operatorname{abs}\left(B_{\left.G_{Q^{r \otimes}\left(1^{2+n}\right)}\right)}\right)$ as, respectively,

$$
\left(\begin{array}{cc}
A & B  \tag{3.19}\\
C & 0
\end{array}\right) \quad \text { and } \quad\left(\begin{array}{ccc}
\alpha_{1} A & \beta_{1} B & 0 \\
\alpha_{2} A & 0 & \beta_{2} B \\
\alpha C & 0 & 0
\end{array}\right)
$$

for some entry-wise nonnegative matrices $A, B, C$. Now if $\beta_{1}$ and $\beta_{2}$ were equal to $\alpha$, then the norm of the right biadjacency matrix would be exactly $\alpha$ times the norm of the left matrix. Since in fact $\beta_{1}=\beta_{2}=1<\alpha$, Eq. (3.16) holds.

Let $\varphi$ be an AND-OR formula of size $n$, in which each AND and each OR gate has fan-in two. Let $r$ be the root of $\varphi$. For each vertex $v$ in $\varphi$, let $P_{v}$ be the span program

$$
P_{v}= \begin{cases}P_{\mathrm{AND}}\left(s_{1}(v), s_{2}(v)\right) & \text { if } g_{v} \text { is an AND gate }  \tag{3.20}\\ P_{\mathrm{OR}}\left(s_{1}(v), s_{2}(v)\right) & \text { if } g_{v} \text { is an OR gate }\end{cases}
$$

Denote by $\alpha(v), \alpha_{j}(v), \beta_{j}(v)$ and $\epsilon(v), \delta(v), \epsilon_{j}(v)$, for $j \in[2]$, the parameters of $P_{v}$ from Definition 3.1.

Let $P_{\varphi}$ be the span program formed by composing the span programs $P_{v}$ according to the structure of $\varphi$ from the root $r$ of $\varphi$ toward the leaves, in an otherwise arbitrary order, using


Figure 3: Examples to illustrate reduced tensor-product composition for $\mathrm{AND}_{2}-\mathrm{OR}_{2}$ formulas. For each formula $\varphi$, of size $n$, the graph $G_{P_{\varphi}}\left(1^{n}\right)$ is displayed. Vertices corresponding to maximal false inputs are so labeled. In (e), the primed variables refer to the span program coefficients for $x_{1} \wedge x_{2}$. Notice in each example that a vertex labeled with input $x \in\{0,1\}^{n}$ is connected exactly to those input bits $j$ with $x_{j}=0$; this is a consequence of Eq. (3.22). Also notice that the graph's structure changes locally as each additional gate is composed onto the end of the formula, e.g., from (d) to (e). However, edge weights change nonlocally. See also Figure 2.
reduced tensor-product composition with respect to the bases of Definition 3.1. Note that $P_{\varphi}$ is strict and monotone. If for each $v, s_{1}(v)$ and $s_{2}(v)$ are set to be the sizes of the respective input subformulas, then wsize $\left(P_{\varphi}\right)=\sqrt{n}$. Figure 3 has several examples.

We can characterize $P_{\varphi}$ in terms of the set of "maximal false" inputs, or minimal zero-certificates, to the formula $\varphi$.

Definition 3.5. An input $x \in B^{n}$ is a maximal false input to $\varphi$ if $\varphi(x)=0$ and flipping any bit of $x$ from 0 to 1 changes the formula evaluation to 1 .

To any maximal false input $x$ corresponds a subtree $T_{x}$ of $\varphi . T_{x}$ is rooted at $r$, and its leaves are exactly the leaves of $\varphi$ corresponding to input bits $k$ with $x_{k}=0$. Note that for each vertex $v \in T_{x}$, if $g_{v}$ is an OR gate, then both of $v$ 's children are also in $T_{x}$; and if $g_{v}$ is an AND gate, then exactly one of $v$ 's children is in $T_{x}$.

Lemma 3.6. Let $U$ be the set of maximal false inputs to $\varphi$. Then $P_{\varphi}$ is given as follows:

- Its inner product space is $V=\mathbf{C}^{U}$.
- Its target vector is

$$
|t\rangle=\sum_{x \in U}\left(\prod_{v \in T_{x}}\left\{\begin{array}{cl}
\alpha_{j(x, v)}(v) & \text { if } g_{v}=\mathrm{AND}_{2}  \tag{3.21}\\
\delta(v) & \text { if } g_{v}=\mathrm{OR}_{2}
\end{array}\right\} \cdot \prod_{v \notin T_{x}}\left\{\begin{array}{cl}
\alpha(v) & \text { if } g_{v}=\mathrm{AND}_{2} \\
\delta(v) & \text { if } g_{v}=\mathrm{OR}_{2}
\end{array}\right\}\right)|x\rangle .
$$

Both products are over internal vertices only (the parameters $\alpha(v), \ldots$ are not defined for leaves). In the first term $j(x, v) \in[2]$ indicates the child of $v$ that is in $T_{x}$.

- Its input vectors are indexed by $I_{k 1}=\{k\}, I_{k 0}=\emptyset$ for $k \in[n]$. Letting $\gamma_{k}$ be the simple path from the $k$ 'th leaf to $r,\left|v_{k}\right\rangle$ is given by

$$
\left|v_{k}\right\rangle=\sum_{x \in U: x_{k}=0}\left(\prod_{v \in \gamma_{k}}\left\{\begin{array}{c}
\beta_{j(k, v)}(v)  \tag{3.22}\\
\epsilon_{j(k, v)}(v)
\end{array}\right\} \cdot \prod_{v \in T_{x} \backslash \gamma_{k}}\left\{\begin{array}{c}
\alpha_{j(x, v)}(v) \\
\delta(v)
\end{array}\right\} \cdot \prod_{v \notin T_{x}}\left\{\begin{array}{c}
\alpha(v) \\
\delta(v)
\end{array}\right\}\right)|x\rangle
$$

Here in each bracketed expression the top term is to be taken if $g_{v}$ is an AND gate, and the bottom term if $g_{v}$ is an $O R$ gate. The indices $j(x, v) \in[2]$ are as in the expression for $|t\rangle$, while $j(k, v) \in[2]$ indicates the child of $v$ that is along the path $\gamma_{k}$. Products are over internal vertices only.

In particular, for $x \in U,\left\langle x \mid v_{k}\right\rangle=0$ if and only if $x_{k}=1$. Thus $|x\rangle /\langle t \mid x\rangle$ is a witness for $f_{P_{\varphi}}(x)=0$.
Proof. Although the expressions are intimidating, the proof is a simple induction. Assume that we have completed reduced tensor-product composition of the span programs $P_{v}$ for vertices $v$ belonging to a subtree $\varphi^{\prime}$ that includes $r$. Then $\varphi^{\prime}$ is also an AND-OR formula. Assume that the characterization of Lemma 3.6 holds for $\varphi^{\prime}$.

Consider adding a new vertex $v$ onto $\varphi^{\prime}$, yielding a formula $\varphi^{\prime \prime}$. We will use primed variables, $U^{\prime}, T_{x}^{\prime}, \gamma_{k}^{\prime}$, to refer to $\varphi^{\prime}$ and double-primed variables for $\varphi^{\prime \prime}$.

The base case of the induction is if $\varphi^{\prime}$ is empty and $\varphi^{\prime \prime}=\{r\}$. Then the claim is a consequence of Definition 3.1. If $g_{r}$ is an AND gate, then $U^{\prime \prime}=\{01,10\}$, while $U^{\prime \prime}=\{00\}$ if $g_{r}$ is an OR gate.

Now assume that the size of $\varphi^{\prime}$ is at least two. By symmetry, assume that the first input of $\varphi^{\prime}$ is replaced by $v$. Then by induction, note that $Z=\left\{\right.$ maximal false inputs $x$ to $\left.\varphi^{\prime} \mid x_{1}=1\right\}$ in Definition 3.3. There are two cases, depending on whether the gate $g_{v}$ is an OR or an AND.

- If the new gate is an OR , then the maximal false inputs of $\varphi^{\prime}$ are in one-to-one correspondence to those of the new formula $\varphi^{\prime \prime}$. By Definition 3.3 the inner product space does not change. Moreover, the target vector is scaled simply by $\delta(v)$. All but the first input vector of $P_{\varphi^{\prime}}$ are scaled by $\delta(v)$. The first input vector is split into two vectors, scaled by $\epsilon_{1}(v)$ and $\epsilon_{2}(v)$. By examination of Eqs. (3.21) and (3.22), the induction assumption is maintained.
- If the new gate is an AND, then a maximal false input $1 x$ of $\varphi^{\prime}$ corresponds to the maximal false input $11 x$ of $\varphi^{\prime \prime}$, while any maximal false input $0 x$ of $\varphi^{\prime}$ splits into the two maximal false inputs $01 x$ and $10 x$ of $\varphi^{\prime \prime}$. Consider first a maximal false input $1 x$. Since $1 x \in Z$, the target and all input vectors are simply scaled by $\alpha(v)$ in this coordinate. This satisfies the induction hypothesis since $v \notin T_{11 x}^{\prime \prime}$.
Next consider a maximal false input $0 x$ to $\varphi^{\prime} ; v \in T_{01 x}^{\prime \prime} \cap T_{10 x}^{\prime \prime}$. By Eq. (3.8), the $0 x$ coordinate of the target vector splits in two, weighted by $\alpha_{1}(v)$ and $\alpha_{2}(v)$, so Eq. (3.21) continues to hold. Similarly, for the unchanged inputs $k$, i.e., inputs with $v \notin \gamma_{k}^{\prime \prime},\binom{\left\langle 01 x \mid v_{k}^{\prime \prime}\right\rangle}{\left\langle 10 x \mid v_{k}^{\prime \prime}\right\rangle}=\binom{\alpha_{1}(v)}{\alpha_{2}(v)}\left\langle 0 x \mid v_{k}^{\prime}\right\rangle$ by Eq. (3.10). This is accounted for by the second term of Eq. (3.22). By Eq. (3.10), $\binom{\left\langle 01 x \mid v_{1}^{\prime \prime}\right\rangle}{\left\langle 10 x \mid v_{1}^{\prime \prime}\right\rangle}=\binom{\beta_{1}(v)}{0}\left\langle 0 x \mid v_{1}^{\prime}\right\rangle$. This shows up in the first term of Eq. (3.22), since $v \in \gamma_{1}^{\prime \prime}$. A similar argument holds for the input vector $\left|v_{2}^{\prime \prime}\right\rangle$.

One straightforward consequence of Lemma 3.6 is that the order in which gates are composed in $P_{\varphi}$ does not matter. In fact, the composition order does not matter for direct-sum, tensor-product or reduced tensor-product composition of arbitrary span programs. Rather than prove this claim, though, in the sequel we will continue to order composition from the root toward the leaves.

Since its construction uses reduced tensor-product composition, $P_{\varphi}$ has similar properties as a canonical span program [Rei09a, Def. 5.1]. A canonical span program has a dimension for every false input, whereas $P_{\varphi}$ has a dimension only for each maximal false input. Unlike a canonical span program, $P_{\varphi}$ 's target vector does not in general have uniform weights. However, if $x \in B^{n}$ has $\varphi(x)=0$, let $x^{\prime}$ be a maximal false input that lies above $x$, i.e., the bitwise AND of $x$ and $x^{\prime}$ is $x$. Then $\left|x^{\prime}\right\rangle /\left\langle t \mid x^{\prime}\right\rangle$ is a witness for $f_{P_{\varphi}}(x)=0$. This property is a main motivation for defining reduced tensor-product span program composition; it does not hold for tensor-product composition.

Unfortunately, a construction based only on reduced tensor-product span program composition will not work for applying Theorem 2.3. There are two problems:

1. First, the number of maximal false inputs can be exponentially large in the size $n$ of the formula. For example, if $\varphi$ is a balanced formula of depth $d$ with alternating levels of $\mathrm{AND}_{2}$ and $\mathrm{OR}_{2}$ gates, starting with AND, then there are $\frac{1}{2} 4^{2^{\lfloor(d-1) / 2\rfloor}}=e^{\Omega(\sqrt{n})}$ maximal false inputs. Since there are only $n$ input vectors, this implies that the maximum degree of a vertex in the graph must be exponentially large.
2. Second, the norm $\left\|\operatorname{abs}\left(A_{G_{P_{\varphi}}}\right)\right\|$ can also be exponentially large. Indeed, $\alpha(v)$ can be as large as $2^{1 / 4}$ when the two input subformulas to $v$ have the same size. Then, for example, a balanced formula of depth $d$ with only $\mathrm{AND}_{2}$ gates leads to a target vector with coefficients each $2^{\frac{1}{4}\left(2^{d}-(2 d+1)\right)}$.

The advantage of reduced tensor-product composition, though, is that because it outputs a strict span program $P_{\varphi}$, the full witness size can be easily bounded in terms of the witness size.

The disadvantages of reduced tensor-product composition are worst for very balanced formulas, while its advantage is most helpful for unbalanced formulas. This suggests that a combining the two techniques might be useful for general general AND-OR formulas.

Before turning to such an approach, though, let us first restate Lemma 3.6 for the case of a maximally unbalanced formula. Lemma 3.7 is a key structural claim behind our proof of Theorem 1.1. Once again, Figure 3 has several examples.

Lemma 3.7. Let $\psi$ be a maximally unbalanced AND-OR formula, i.e., the depth of the formula equals the number of gates $J$. Index the inputs in order from farthest to closest to the root (see Figure 4). Let $T_{\mathrm{AND}}=\{j \in[J+1]: j$ is an input to an AND gate $\}$, and $T_{\mathrm{OR}}=[J+1] \backslash T_{\mathrm{AND}}$.

Then the maximal false inputs $U$ to $\psi$ are in one-to-one correspondence to the elements of $T_{F}=T_{\mathrm{AND}} \cup\{1\}$. To $j \in T_{F}$ corresponds the maximal false input $x^{j} \in B^{J+1}$ given by

$$
x_{k}^{j}= \begin{cases}0 & \text { if } k=j, \text { or if } k>j \text { and } k \in T_{\mathrm{OR}}  \tag{3.23}\\ 1 & \text { if } k<j, \text { or if } k>j \text { and } k \in T_{\mathrm{AND}}\end{cases}
$$

The reduced-tensor-product-composed span program $P_{\psi}$ is given as follows:

- Its inner product space is $V=\mathbf{C}^{U}$.
- Its target vector is

$$
|t\rangle=\sum_{j \in T_{F}}\left(\prod_{v \in \gamma_{j}}\left\{\begin{array}{c}
\alpha_{\iota(j, v)}(v)  \tag{3.24}\\
\delta(v)
\end{array}\right\} \cdot \prod_{v \notin \gamma_{j}}\left\{\begin{array}{c}
\alpha(v) \\
\delta(v)
\end{array}\right\}\right)\left|x^{j}\right\rangle .
$$

Here, in the first term $\iota(j, v) \in[2]$ indicates the child of $v$ that is in $\gamma_{j}$, the simple path from the jth vertex to the root. In each bracketed expression the top term is to be taken if $g_{v}$ is an AND gate, and the bottom term if $g_{v}$ is an OR gate. The products are over internal vertices $v$.

- Its input vectors are indexed by $I_{k 1}=\{k\}, I_{k 0}=\emptyset$ for $k \in[J+1]$. For $k \in T_{\mathrm{AND}} \cup\left[\min T_{\mathrm{AND}}\right]$, $\left|v_{k}\right\rangle$ is given by

$$
\left|v_{k}\right\rangle=\prod_{v \in \gamma_{k}}\left\{\begin{array}{c}
\beta_{\iota(k, v)}(v)  \tag{3.25}\\
\epsilon_{\iota(k, v)}(v)
\end{array}\right\} \cdot \prod_{v \notin \gamma_{k}}\left\{\begin{array}{c}
\alpha(v) \\
\delta(v)
\end{array}\right\}\left|x^{\kappa}\right\rangle
$$

where $\kappa=k$ if $k \in T_{\text {AND }}$ and $\kappa=1$ otherwise. For $k \notin T_{\mathrm{AND}} \cup\left[\min T_{\mathrm{AND}}\right]$,

$$
\left|v_{k}\right\rangle=\sum_{\substack{j \in T_{F}  \tag{3.26}\\
j<k}}\left(\prod_{v \in \gamma_{k}}\left\{\begin{array}{c}
\beta_{\iota(k, v)}(v) \\
\epsilon_{\iota(k, v)}(v)
\end{array}\right\} \cdot \prod_{v \in \gamma_{j} \backslash \gamma_{k}}\left\{\begin{array}{c}
\alpha_{\iota(j, v)}(v) \\
\delta(v)
\end{array}\right\} \cdot \prod_{v \notin \gamma_{j}}\left\{\begin{array}{c}
\alpha(v) \\
\delta(v)
\end{array}\right\}\right)\left|x^{j}\right\rangle .
$$

The same conventions are understood as in the expression for $|t\rangle$.
In particular, for $x \in U,|x\rangle /\langle t \mid x\rangle$ is a witness for $f_{P_{\psi}}(x)=0$.
Proof. The lemma follows from Lemma 3.6. The main simplification to make is that the tree $T_{x^{j}}=\gamma_{j} \cup\left\{k \in T_{\mathrm{OR}}: k>j\right\}$, so $T_{x^{j}}$ and $\gamma_{j}$ agree on internal vertices. Then Eq. (3.22) simplifies to

$$
\left|v_{k}\right\rangle=\sum_{j \in T_{F}: x_{k}^{j}=0}\left(\prod_{v \in \gamma_{k}}\left\{\begin{array}{c}
\beta_{\iota(k, v)}(v)  \tag{3.27}\\
\epsilon_{\iota(k, v)}(v)
\end{array}\right\} \cdot \prod_{v \in \gamma_{j} \backslash \gamma_{k}}\left\{\begin{array}{c}
\alpha_{\iota(j, v)}(v) \\
\delta(v)
\end{array}\right\} \cdot \prod_{v \notin \gamma_{j}}\left\{\begin{array}{c}
\alpha(v) \\
\delta(v)
\end{array}\right\}\right)\left|x^{j}\right\rangle
$$

This further simplifies to Eq. (3.25) when $k \in T_{\mathrm{AND}} \cup\left[\min T_{\mathrm{AND}}\right]$ because then the sum over $j \in T_{F}$ with $x_{k}^{j}=0$ has only a single term, either $k$ itself or 1 if $k<\min T_{\text {AND }}$. It simplifies to Eq. (3.26) when $k \notin T_{\mathrm{AND}} \cup\left[\min T_{\mathrm{AND}}\right]$ because then $x_{k}^{j}=0$ for $j \in T_{F}$ exactly when $j<k$.

### 3.3 AND-OR span program construction using hybrid composition

Let $\varphi$ be an AND-OR formula in which every gate has fan-in two, possibly after expanding higher fan-in gates. We now construct the span program $P_{\varphi}$ on which the algorithm in Theorem 1.1 is based. For each internal vertex $v$ of $\varphi$, fix the parameters $s_{1}(v)$ and $s_{2}(v)$ to the sizes of the respective input subformulas. Assume without loss of generality that $s_{1}(v) \geq s_{2}(v)$ always.

To begin the construction, we mark certain edges of the formula. Marked edges are termed "checkpoints," because they will serve to cut off reduced tensor-product span program compositions and the associated exponential growths in norm and degree. There are two steps to placing the checkpoints:

1. For every internal vertex $v$, i.e., a vertex that is not a leaf, mark the edge to the smaller of its two input subformulas. Break ties arbitrarily.
After this step, every gate has exactly one unmarked input edge, so the marked edges divide $\varphi$ into a set of paths. Let

$$
\begin{equation*}
\mathcal{S}=\{\text { vertices } v: \text { for one of the paths, } v \text { is the endpoint closer to the root } r\} . \tag{3.28}
\end{equation*}
$$

Certainly the root $r$ is itself in $\mathcal{S}$.
2. In the second step, place more checkpoints to split up paths that are too long. For each path, apply the following rule:
Starting at the far end of the path and moving toward the root, keep track of the product

$$
\prod_{v}\left\{\begin{array}{ll}
\alpha(v) & \text { if } g_{v}=\mathrm{AND}_{2}  \tag{3.29}\\
\epsilon(v) & \text { if } g_{v}=\mathrm{OR}_{2}
\end{array}\right\}=\prod_{v}\left(\frac{\sqrt{s_{1}(v)}+\sqrt{s_{2}(v)}}{\sqrt{s_{1}(v)+s_{2}(v)}}\right)^{1 / 2}
$$

for internal vertices $v$ along the path. Note that $\alpha(v), \epsilon(v) \leq 2^{1 / 4} \approx 1.19$. After adding a vertex that makes the product exceed $\sqrt{e} \approx 1.65$, split the path by adding a checkpoint.
After finishing this step, there will be no paths with the above product more than $2^{1 / 4} \sqrt{e}$, and for all paths except possibly those ending at a vertex in $\mathcal{S}$, the product will be at least $\sqrt{e}$.

We remark that our analysis in Section 3.4 is not overly sensitive to modifying these rules, for example by using a different constant greater than one instead of $\sqrt{e}$ in the second step.

Based on the checkpointed formula, we can now construct the composed span program $P_{\varphi}$ :

1. First, for each path $\xi$, compose the span programs $P_{v}$ from Eq. (3.20) along the path, using reduced tensor-product composition, to obtain a span program $P_{\xi}$.
2. Next, apply direct-sum composition to compose the $P_{\xi}$ span programs constructed in the first step across the checkpointed edges.

Let $P_{\varphi}$ be the resulting span program. Figure 2 shows an example. Notice that when reduced tensor-product and direct-sum span program composition are both used, their relative order matters. In constructing $P_{\varphi}$, all direct-sum composition comes last.

### 3.4 Analysis of $P_{\varphi}$

Proof of Theorem 1.1. From Claim 3.2 and the span program composition results [Rei09a, Theorem 4.3, Prop. 4.7], we obtain:
Lemma 3.8. $P_{\varphi}$ computes $f_{P_{\varphi}}=\varphi$ and $\operatorname{wsize}\left(P_{\varphi}\right)=\sqrt{n}$.
In order to apply Theorem 2.3, the basic idea is to treat the gates along each checkpointed path $\xi$ as a single grouped gate, and to analyze the direct-sum composition of these grouped gates as in the proof of [Rei09b, Theorem 1.11]. There are two steps to this analysis.

1. First, we argue that $\left\|\operatorname{abs}\left(A_{G_{P_{\varphi}}}\right)\right\|=O(1)$ and $G_{P_{\varphi}}$ has maximum degree $O(\sqrt{n})$.
2. Second, to bound the full witness size we apply Lemma 2.5 to the direct-sum composition of the grouped gates, similarly to [Rei09b, Lemmas 3.3, 4.4].

Let us establish a bound on $\left\|A_{G_{P_{\varphi}}}\right\|$. Since $P_{\varphi}$ has nonnegative entries, $\operatorname{abs}\left(A_{G_{P_{\varphi}}}\right)=A_{G_{P_{\varphi}}}$.
Lemma 3.9. The norm of the adjacency matrix $A_{G_{P_{\varphi}}}$ is at most $2(2 \sqrt{2} e+1)=O(1)$.
Proof. By Lemma 2.6, it suffices to bound the norm of $A_{G_{P_{\xi}}}$ for any checkpointed path $\xi$.
By definition, if a path $\xi$ involves $J$ internal vertices,

$$
A_{G_{P_{\xi}}}=\left(\begin{array}{ccc}
0 & 0 & B  \tag{3.30}\\
0 & 0 & C \\
B^{\dagger} & C^{\dagger} & 0
\end{array}\right)
$$

where $B$ is the biadjacency matrix $B_{G_{P_{\xi}\left(1^{J+1}\right)}}$ and $C$ is a column of zeros followed by an identity matrix. Therefore we bound

$$
\left\|A_{G_{P_{\xi}}}\right\| \leq\left\|\left(\begin{array}{cc}
0 & B  \tag{3.31}\\
B^{\dagger} & 0
\end{array}\right)\right\|+\left\|\left(\begin{array}{cc}
0 & C \\
C^{\dagger} & 0
\end{array}\right)\right\|=\|B\|^{2}+1 .
$$

Now from Lemma 3.4 and the checkpoint rule Eq. (3.29), $\|B\|^{2} \leq \sqrt{2} e\|(1 \quad 1)\|^{2}=2 \sqrt{2} e$.
For a path $\xi$, let $\psi(\xi)$ be the AND-OR formula that is the composition of the gates along the path $\xi$. $\psi(\xi)$ is maximally unbalanced, so Lemma 3.7 applies to $P_{\xi}=P_{\psi(\xi)}$. In particular, this allows us to bound the maximum degree of a vertex in $G_{P_{\xi}}$, and to bound the lengths of witnesses to $f_{P_{\psi(\xi)}}(x)=0$ :
Corollary 3.10. The maximum degree of a vertex in $G_{P_{\varphi}}$ is $O(\sqrt{n})$.
Proof. Since $P_{\varphi}$ is constructed by direct-sum composition of span programs $P_{\xi}$, the maximum degree of a vertex in $G_{P_{\varphi}}$ is at most twice the maximum degree of a vertex in a graph $G_{P_{\xi}}$, which is at most twice the number of vertices in $G_{P_{\xi}}$.

Fix a path $\xi$ with $J$ gates. By Lemma 3.7, the number of vertices in $G_{P_{\xi}}\left(1^{J+1}\right)$ is $1+\left|T_{F}\right|+$ $(J+1) \leq 2 J+3$. Now $\prod_{v \in \xi} \alpha(v) \leq 2^{1 / 4} \sqrt{e}$, but each $\alpha(v)$ satisfies

$$
\begin{equation*}
\alpha(v)^{2}>\frac{1+\frac{1}{\sqrt{n}}}{\sqrt{1+\frac{1}{n}}}>1+\frac{1}{2 \sqrt{n}}, \tag{3.32}
\end{equation*}
$$

where in the first inequality we used $s_{2}(v) \geq 1$ and $s_{1}(v)<n$. Hence $J=O(\sqrt{n})$.
Corollary 3.11. For a checkpointed path $\xi$, let $\psi=\psi(\xi)$, and recall from Lemma 3.7 the definitions of $T_{F}=T_{\mathrm{AND}} \cup\{1\}$ and of the maximal false inputs $x^{j}$ for $j \in T_{F}$. Let $m$ be the size of the subformula of $\varphi$ rooted at the root of $\xi$.

Then for any $j \in T_{F},\left|w_{j}^{\prime}\right\rangle=\left|x^{j}\right\rangle /\left\langle t \mid x^{j}\right\rangle$ is a witness for $f_{P_{\xi}}\left(x^{j}\right)=0$, with

$$
\|\left|w_{j}^{\prime}\right\rangle \|^{2} \leq \begin{cases}\sqrt{2} e & \text { if } j=1  \tag{3.33}\\ \sqrt{2} e \sqrt{m} & \text { otherwise }\end{cases}
$$



Figure 4: A typical checkpointed path $\xi$ in $\varphi$ on $J=6$ vertices. The root $r$ is to the left. Input subformulas are indexed from right to left; of 1 and 2 , the larger subformula comes first.

Proof. The coefficient $\left\langle t \mid x^{j}\right\rangle$ has been computed in Eq. (3.24). Since $\alpha(v) \geq 1$ and $\delta(v)=1$ always, we need to place an upper bound on $1 /\left|\left\langle t \mid x^{j}\right\rangle\right|^{2} \leq \prod_{v \in \gamma_{j}} \alpha_{\iota(j, v)}(v)^{-2}$.

Note that

$$
\begin{equation*}
\frac{1}{\alpha_{1}(v)^{2}}=\sqrt{\frac{s_{1}(v)+s_{2}(v)}{s_{1}(v)}}=\sqrt{1+r(v)}, \tag{3.34}
\end{equation*}
$$

where $r(v)=s_{2}(v) / s_{1}(v) \leq 1$. Since $\alpha(v)^{2}=(1+\sqrt{r(v)}) / \sqrt{1+r(v)}$, therefore $\frac{1}{\alpha_{1}(v)^{2}} \leq \alpha(v)^{2}$. On the other hand, the best bound we can place on $\frac{1}{\alpha_{2}(v)^{2}}$ is, since $s_{2}(v) \geq 1$,

$$
\begin{equation*}
\frac{1}{\alpha_{2}(v)^{2}}=\sqrt{\frac{s_{1}(v)+s_{2}(v)}{s_{2}(v)}} \leq \sqrt{s_{1}(v)+s_{2}(v)} . \tag{3.35}
\end{equation*}
$$

When $j=1, \iota(j, v)=1$ for every $v$, so we obtain the claimed bound

$$
\begin{equation*}
\|\left|w_{j}^{\prime}\right\rangle \|^{2} \leq \prod_{v \in \xi} \alpha(v)^{2} \leq \sqrt{2} e \tag{3.36}
\end{equation*}
$$

For $j>1$, all but one of the indices $\iota(j, v)$ will be 1 , and for the last internal vertex $v$ on $\gamma_{j}, \iota(j, v)$ will be 2. It follows that $\|\left|w_{j}^{\prime}\right\rangle \|^{2} \leq\left(\prod_{v \in \xi} \alpha(v)^{2}\right) \sqrt{m} \leq \sqrt{2} e \sqrt{m}$.

Let $\varphi^{\prime}$ be a formula whose internal vertices are the checkpointed paths in $\varphi$ and whose leaves correspond to the inputs of $\varphi$. The formula $\varphi^{\prime}$ is not an AND-OR formula; its gate at an internal vertex $v^{\prime}$ is the composition of the gates in the corresponding path $\xi\left(v^{\prime}\right)$ in $\varphi, g_{v^{\prime}}^{\prime}=f_{P_{\xi\left(v^{\prime}\right)}}$. However, $\varphi^{\prime}$ and $\varphi$ compute the same boolean function. Let $r^{\prime}$ be the root of $\varphi^{\prime}$, corresponding to the path in $\varphi$ that includes $r$. For each vertex $v^{\prime}$ of $\varphi^{\prime}$, let $\varphi_{v^{\prime}}^{\prime}$ be the subformula rooted at $v^{\prime}$, and let $P_{\varphi_{v^{\prime}}^{\prime}}$ be the span program obtained by direct-sum composition of the span programs $P_{\xi\left(w^{\prime}\right)}$ for all vertices $w^{\prime}$ in $\varphi_{v^{\prime}}^{\prime}$. Then $P_{\varphi}=P_{\varphi_{r^{\prime}}^{\prime}}$. Let $s_{v^{\prime}}$ be the number of inputs to $\varphi_{v^{\prime}}^{\prime}$, so wsize $\left(P_{\varphi_{v^{\prime}}^{\prime}}^{\prime}\right)=\sqrt{s_{v^{\prime}}}$.

Call an internal vertex $v^{\prime}$ of $\varphi^{\prime}$ small if the corresponding path $\xi\left(v^{\prime}\right)$ has an endpoint in the set $\mathcal{S}$ from Eq. (3.28). Aside from the root $r^{\prime}$, every small vertex $v^{\prime}$ has subformula size $s_{v^{\prime}}$ no larger than the size of its sibling subformula in $\varphi$; hence the name.

Lemma 3.12. Let $v^{\prime}$ be an internal vertex of $\varphi^{\prime}$, with children $c_{1}^{\prime}, c_{2}^{\prime}, \ldots, c_{J+1}^{\prime}$, sorted in decreasing order of their distances from $r^{\prime}$, and with $s_{c_{1}^{\prime}} \geq s_{c_{2}^{\prime}}$ (see Figure 4). Then for $j \geq 2, s_{c_{j}^{\prime}} \leq s_{v^{\prime}} / 2$. If $v^{\prime}$ is not a small vertex, then $s_{c_{1}^{\prime}} \leq s_{v^{\prime}}-\frac{1}{\sqrt{2}} \sqrt{s_{v^{\prime}}}$.

Proof. In placing checkpoints, at every vertex $v$ the edge to the smaller input subformula is marked. Thus the size of that subformula can be at most half the size of $\varphi_{v}$. However the larger subformula
can have size up to $s_{v}-1$. Assuming that $v^{\prime}$ is not small, we will argue that in the second step of placing checkpoints, paths are allowed to grow long enough toward the root that the size increases significantly.

Say that in the second step of placing checkpoints we begin with a single vertex whose input subformulas have size $s$ and $r_{1} s$, where $r_{1} \in[1 / s, 1]$. The next vertex toward the root will have another input subformula whose size is, say, $r_{2}\left(s+r_{1} s\right)$, where $r_{2} \in\left[1 /\left(s+r_{1} s\right), 1\right]$. Similarly the $j$ th vertex towards the root $r$ will have a smaller input subformula of size $r_{j} s \prod_{k<j}\left(1+r_{k}\right)$, so its two subformulas have total size $s \prod_{k \leq j}\left(1+r_{k}\right)$. Here, $r_{j} \leq 1$ and $1 / r_{j} \leq s \prod_{k<j}\left(1+r_{k}\right)$.

Now it is possible that this path gets all the way to $\mathcal{S}$ without a new checkpoint being placed. Then this path corresponds to a small vertex, and we can make no strong claim about how much the size increases along the path. For example, if the path starts at the root $r$, then no further checkpoints can be placed, so we can only bound $s_{c_{1}^{\prime}} \leq n-1$.

Assume that a checkpoint is placed to terminate this path. Then we have $\prod_{j} \frac{1+\sqrt{r_{j}}}{\sqrt{1+r_{j}}} \geq e$. We want to lower bound $\prod_{j}\left(1+r_{j}\right)$. Letting $x_{j}=\log \frac{1+\sqrt{r_{j}}}{\sqrt{1+r_{j}}} \in\left(0, \frac{1}{2} \log 2\right], \log \left(1+r_{j}\right)=\log \frac{2}{1+e^{x_{j}} \sqrt{2-e^{2 x_{j}}}}$. Applying Jensen's inequality gives $\prod_{j}\left(1+r_{j}\right) \geq e^{1 / J}>1+1 / J$, where $J$ is the number of vertices along the path.

We have $s_{v^{\prime}}=s_{c_{1}^{\prime}} \prod_{j}\left(1+r_{j}\right) \geq 2$. If $J \geq \frac{1}{\sqrt{2}} \sqrt{s_{v^{\prime}}}$, then since at each step toward the root the size must increase by at least one, $s_{c_{1}^{\prime}} \leq s_{v^{\prime}}-J \leq s_{v^{\prime}}-\frac{1}{\sqrt{2}} \sqrt{s_{v^{\prime}}}$, as claimed. If $J<\frac{1}{\sqrt{2}} \sqrt{s_{v^{\prime}}}$, then using $s_{c_{1}^{\prime}}<s_{v^{\prime}} /(1+1 / J)$ again gives $s_{c_{1}^{\prime}}<s_{v^{\prime}}-\frac{1}{\sqrt{2}} \sqrt{s_{v^{\prime}}}$.

In the maximally unbalanced formula, $s_{c_{1}^{\prime}}=s_{v^{\prime}}-O\left(\sqrt{s_{v^{\prime}}}\right)$, so aside from improving the constant $\frac{1}{\sqrt{2}}$, the bound on $s_{c_{1}^{\prime}}$ in Lemma 3.12 is tight.
Lemma 3.13. Let $v^{\prime}$ be an internal vertex of $\varphi^{\prime}$, and let $m=s_{v^{\prime}}$. Then

$$
\frac{\operatorname{fwsize}\left(P_{\varphi_{v^{\prime}}}\right)}{\sqrt{m}} \leq \begin{cases}\kappa \log m & \text { if } v^{\prime} \text { is not small }  \tag{3.37}\\ \kappa \log m+\frac{\lambda}{\sqrt{m}} & \text { if } v^{\prime} \text { is small }\end{cases}
$$

where $\lambda=\sqrt{2} e \approx 3.84$ and $\kappa=\left(1+\frac{1}{\sqrt{2}}\right) \lambda / \log 2 \approx 9.47$.
Proof. The proof is by induction in the maximum distance from $v^{\prime}$ to a leaf in $\varphi^{\prime}$. The base case is if all of $v^{\prime}$ 's inputs are leaves. This case can be handled simultaneously to the induction step by letting $S=\emptyset$, so $\sigma=1$, in the applications of Lemma 2.5 below.

Take $v^{\prime}$ an internal vertex, with children $c_{1}^{\prime}, c_{2}^{\prime}, \ldots, c_{J+1}^{\prime}$, sorted in decreasing order of their distances from $r^{\prime}$, and with $s_{c_{1}^{\prime}} \geq s_{c_{2}^{\prime}}$. Note that each vertex $c_{j}^{\prime}$ with $j \geq 2$ must be either small or a leaf. However, $c_{1}^{\prime}$ is not a small vertex. Let $\xi=\xi\left(v^{\prime}\right)$ be the corresponding checkpointed path in $\varphi$. Consider an input $x \in B^{m}$. We will use Lemma 2.5 in combination with Lemma 3.12 and Corollary 3.11.

If $\varphi_{v^{\prime}}^{\prime}(x)=1$, then apply Lemma 2.5 with $S=\left\{j \in[J+1]: c_{j}^{\prime}\right.$ is not a leaf $\}$. Since $P_{\xi\left(v^{\prime}\right)}$ is a strict span program, Eq. (2.10) gives

$$
\begin{equation*}
\frac{\operatorname{fwsize}\left(P_{\varphi_{v^{\prime}}^{\prime}}, x\right)}{\sqrt{m}} \leq \sigma+\frac{1}{\sqrt{m}} \tag{3.38}
\end{equation*}
$$

where $\sigma=\max _{j \in S}$ fwsize $\left(P_{\varphi_{c_{j}^{\prime}}^{\prime}}\right) / \sqrt{\bar{s}_{c_{j}^{\prime}}}$. Now by Lemma 3.12, for $j \in S$ with $j \geq 2, s_{c_{j}^{\prime}} \leq m / 2$ and by induction

$$
\begin{equation*}
\operatorname{fwsize}\left(P_{\varphi_{c_{j}^{\prime}}^{\prime}}\right) / \sqrt{s_{c_{j}^{\prime}}} \leq \kappa \log s_{c_{j}^{\prime}}+\lambda / \sqrt{s_{c_{j}^{\prime}}} . \tag{3.39}
\end{equation*}
$$

Since $\kappa>\lambda /(2 \sqrt{2})$, the right-hand side is an increasing function of $s_{c_{j}^{\prime}}$, so may be bounded by $\kappa \log (m / 2)+\lambda / \sqrt{m / 2}$. For the case $j=1$, by induction fwsize $\left(P_{\varphi_{c_{1}^{\prime}}}\right) / \sqrt{s_{c_{1}^{\prime}}} \leq \kappa \log s_{c_{1}^{\prime}}$.

If $v^{\prime}$ is not small, then Lemma 3.12 gives $s_{c_{1}^{\prime}} \leq m-\sqrt{m / 2}$, and therefore

$$
\begin{equation*}
\sigma \leq \max \left\{\kappa \log (m-\sqrt{m / 2}), \kappa \log (m / 2)+\frac{\lambda}{\sqrt{m / 2}}\right\} \tag{3.40}
\end{equation*}
$$

The first term, bounded by $\kappa(\log m-1 / \sqrt{2 m})$, is at most $\kappa \log m-1 / \sqrt{m}$, provided that $\kappa \geq \sqrt{2}$. The second term satisfies the same bound provided $\kappa \geq\left(\lambda+\frac{1}{\sqrt{2}}\right) / \log 2$.

On the other hand, if $v^{\prime}$ is small, then we can only be sure that $s_{c_{1}^{\prime}} \leq s_{v^{\prime}}-1=m-1$. This is all right, because Eq. (3.37) allows for more slack in this case. In particular, it is enough to show that $\kappa \log (m / 2)+\lambda / \sqrt{m / 2} \leq \kappa \log m$, which indeed holds for $\kappa \geq \lambda / \log 2$.

Assume now that $\varphi_{v^{\prime}}^{\prime}(x)=0$. Let $x^{\prime}$ be the input to $g_{v^{\prime}}^{\prime}$, i.e., the evaluations of $v^{\prime}$ 's input subformulas on $x$. Let $y^{\prime}$ be any minimal false input (Definition 3.5) to $g_{v^{\prime}}^{\prime}$ that lies above $x^{\prime}$ (meaning that for every $k, y_{k}^{\prime} \geq x_{k}^{\prime}$ ). In the notation of Lemma 3.7, $y^{\prime}=x^{j}$ for some $j \in T_{F}$, and $\left|w_{j}^{\prime}\right\rangle=\left|x^{j}\right\rangle /\left\langle t \mid x^{j}\right\rangle$ is a witness for $f_{P_{\xi\left(v^{\prime}\right)}}\left(x^{\prime}\right)=0$. By Eq. (2.11), then,

$$
\begin{align*}
& \frac{\operatorname{fwsize}\left(P_{\varphi_{v^{\prime}}^{\prime}}, x\right)}{\sqrt{m}} \leq \sigma\left(\bar{x}^{\prime},\left|w_{j}^{\prime}\right\rangle\right)+\frac{\|\left|w_{j}^{\prime}\right\rangle \|^{2}}{\sqrt{m}} \\
& \quad \text { where } \sigma\left(\bar{x}^{\prime},\left|w_{j}^{\prime}\right\rangle\right)=\max _{\substack{k \in: \\
x_{k}^{\prime}=0 \text { and }\left\langle v_{k} \mid w_{j}^{\prime}\right\rangle \neq 0}} \frac{\operatorname{fwsize}\left(P_{\varphi_{c_{k}^{\prime}}^{\prime}}\right.}{\sqrt{\mathcal{S}_{k}^{\prime}}} . \tag{3.41}
\end{align*}
$$

There are two cases to consider, $j>1$ and $j=1$.
First, if $j>1$, then by Lemma 3.7, $\left\langle v_{k} \mid w_{j}^{\prime}\right\rangle=0$ for all $k<j$. Therefore, $\sigma=\sigma\left(\bar{x}^{\prime},\left|w_{j}^{\prime}\right\rangle\right)$ will only maximize over small vertices, $c_{k}^{\prime}$ with $k>1$. By induction, $\sigma \leq \kappa \log (m / 2)+\lambda / \sqrt{m / 2}$. On the other hand, Corollary 3.11 gives only $\|\left|w_{j}^{\prime}\right\rangle \|^{2} \leq \lambda \sqrt{m}$. Putting these bounds together,

$$
\begin{align*}
\frac{\operatorname{fwsize}\left(P_{\varphi_{v}^{\prime}}, x\right)}{\sqrt{m}} & \leq \kappa \log (m / 2)+\lambda / \sqrt{m / 2}+\lambda  \tag{3.42}\\
& \leq \kappa \log m
\end{align*}
$$

provided that $\kappa \geq 2 \lambda / \log 2$. The above bound holds whether or not $v^{\prime}$ is small.
Next consider the case $j=1$. Then we will have a worse bound on $\sigma$, but Corollary 3.11 gives $\|\left|w_{j}^{\prime}\right\rangle \|^{2} \leq \lambda$. If $j=1$ and $v^{\prime}$ is not small, then Lemma 3.12 gives $s_{c_{1}^{\prime}} \leq m-\sqrt{m / 2}$. Using the induction hypothesis, we obtain

$$
\begin{equation*}
\frac{\operatorname{fwsize}\left(P_{\varphi_{v^{\prime}}^{\prime}}, x\right)}{\sqrt{m}} \leq \max \left\{\kappa \log (m-\sqrt{m / 2}), \kappa \log (m / 2)+\frac{\lambda}{\sqrt{m / 2}}\right\}+\frac{\lambda}{\sqrt{m}} \tag{3.43}
\end{equation*}
$$

This is indeed at most $\kappa \log m$ provided that $\kappa \geq\left(1+\frac{1}{\sqrt{2}}\right) \lambda / \log 2$.

If $j=1$ and $v^{\prime}$ is small, then we can only be sure that $s_{c_{1}^{\prime}} \leq m-1$. However, just as before, this is enough, since then $\kappa \log s_{c_{1}^{\prime}}<\kappa \log m$, while $\kappa \log (m / 2)+\lambda / \sqrt{m / 2} \leq \kappa \log m$ for $\kappa \geq \lambda / \log 2$.

Substituting $v^{\prime}=r^{\prime}$ into Lemma 3.13 gives fwsize $\left(P_{\varphi_{r^{\prime}}^{\prime}}\right)=$ fwsize $\left(P_{\varphi}\right)=O(\sqrt{n} \log n)$. Lemma 3.9 bounds the norm of $A_{G_{P_{\varphi}}}$ and Corollary 3.10 bounds the maximum degree of a vertex. Therefore, Theorem 2.3 applies, completing the proof of Theorem 1.1.

## 4 Open problems

It is possible that our analysis here is loose. Lemma 2.5 is conservative, as it uses the worst input ratio fwsize $\left(P^{j}\right) / \operatorname{wsize}\left(P^{j}\right)$. Potentially a more careful analysis of the full witness size across span programs $P_{\xi}$ could improve our upper bound. However, we do not believe that the full witness size of $P_{\varphi}$ is $O(\sqrt{n})$. More ideas seem to be needed to find an $O(\sqrt{n})$-query quantum algorithm for evaluating arbitrary size- $n$ AND-OR formulas, particularly if we wish the algorithm also to have poly-logarithmic time overhead after preprocessing. One very simple idea would be to start with different span programs for AND and OR. By [Rei09a, Lemma 4.12], there are many possibilities; the same invertible linear transformation can be applied to the target and all input vectors without changing the witness size. We have unsuccessfully investigated choices to see whether reduced tensor-product composition alone would be sufficient to build an optimal span program for which $\left\|\operatorname{abs}\left(A_{G_{P_{\varphi}}}\right)\right\|=O(1)$, but have not studied other choices under hybrid span program composition.

Some preprocessing will always be necessary if the formula $\varphi$ is presented poorly to the algorithm. However, we would like the preprocessing to correspond to as natural a presentation of $\varphi$ as possible, and there is clearly more work to be done.

Theorem 1.1 and [Rei09b, Theorem 1.11] can likely both be generalized to apply to formulas over a larger class of gates. In particular, the idea of combining direct-sum and tensor-product span program composition might be useful in other contexts. It leads to a dramatic speedup on AND-OR formulas that have not been rebalanced. Since rebalancing and its effect on the adversary bounds is poorly understood for formulas over larger gate sets, this seems like promising technology to apply. We have specialized to AND-OR formulas because they form an important class of formulas, and because the span programs for the AND and OR gates are especially simple. Further applications will require studying other particular span programs more carefully.

Although our algorithm can be analyzed without the span program formalism, it would not have been discovered without a certain familiarity with span programs. Will span programs be useful for finding quantum algorithms for problems beyond formula evaluation? One approach is to try to solve the semi-definite program (SDP) for the optimal span program witness size [Rei09a]. This SDP is exponentially large, so presumably is difficult to solve exactly, but for special cases perhaps it can be solved to within a constant factor of the optimum. Even if so, the span program this gives will typically correspond to a graph with large norm and high degree. To turn this into an algorithm that is time efficient, and not just query efficient, it will be useful to find techniques for breaking up large span programs, perhaps by using tailored semi-definite programs. It is known how to convert a span program into a strict span program without affecting the witness size [Rei09a, Prop. 4.10]. However we do not know of any techniques for converting a span program with low witness size into a span program with low full witness size that additionally has smaller norm and lower degree.

## Acknowledgements

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

## References

$\left[\mathrm{ACR}^{+} 07\right]$ Andris Ambainis, Andrew M. Childs, Ben W. Reichardt, Robert Špalek, and Shengyu Zhang. Any AND-OR formula of size $N$ can be evaluated in time $N^{1 / 2+o(1)}$ on a quantum computer. In Proc. 48th IEEE FOCS, pages 363-372, 2007.
[Amb07] Andris Ambainis. A nearly optimal discrete query quantum algorithm for evaluating NAND formulas. 2007, arXiv:0704.3628 [quant-ph].
[BB94] Maria Luisa Bonet and Samuel R. Buss. Size-depth tradeoffs for Boolean formulae. Information Processing Letters, 49(3):151-155, 1994.
[BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. SIAM J. Comput., 26(5):1510-1523, 1997, arXiv:quant-ph/9701001.
[BCE91] Nader H. Bshouty, Richard Cleve, and Wayne Eberly. Size-depth tradeoffs for algebraic formulae. In Proc. 32nd IEEE FOCS, pages 334-341, 1991.
[BCW98] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In Proc. 30th ACM STOC, pages 63-68, 1998, arXiv:quant-ph/9802040.
[BS04] Howard Barnum and Michael Saks. 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.
[FGG07] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum algorithm for the Hamiltonian NAND tree. 2007, arXiv:quant-ph/0702144.
[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In Proc. 28th ACM STOC, pages 212-219, 1996, arXiv:quant-ph/9605043.
[Gro02] Lov K. Grover. Tradeoffs in the quantum search algorithm. 2002, arXiv:quant-ph/0201152.
[HLŠ07] Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In Proc. 39th ACM STOC, pages 526-535, 2007, arXiv:quant-ph/0611054.
[HMW03] Peter Høyer, Michele Mosca, and Ronald de Wolf. Quantum search on bounded-error inputs. In Proc. 30th ICALP, pages 291-299, 2003, arXiv:quant-ph/0304052. LNCS 2719.
[HW91] Rafi Heiman and Avi Wigderson. Randomized vs. deterministic decision tree complexity for read-once boolean functions. Computational Complexity, 1(4):311-329, 1991. Earlier version in Structure in Complexity Theory '91.
[JKS03] T. S. Jayram, Ravi Kumar, and D. Sivakumar. Two applications of information complexity. In Proc. 35th ACM STOC, pages 673-682, 2003.
[KW93] Mauricio Karchmer and Avi Wigderson. On span programs. In Proc. 8th IEEE Symp. Structure in Complexity Theory, pages 102-111, 1993.
[LLS06] Sophie Laplante, Troy Lee, and Mario Szegedy. The quantum adversary method and classical formula size lower bounds. Computational Complexity, 15:163-196, 2006, arXiv:quant-ph/0501057. Earlier version in Complexity'05.
[NC00] Michael A. Nielsen and Isaac L. Chuang. Quantum computation and quantum information. Cambridge University Press, Cambridge, 2000.
[Rei09a] Ben W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every boolean function. 2009, arXiv:0904.2759 [quant-ph].
[Rei09b] Ben W. Reichardt. Span-program-based quantum algorithm for evaluating unbalanced formulas. 2009, arXiv:0907.1622 [quant-ph].
[RŠ08] Ben W. Reichardt and Robert Špalek. Span-program-based quantum algorithm for evaluating formulas. In Proc. 40th ACM STOC, pages 103-112, 2008, arXiv:0710. 2630 [quant-ph].
[SW86] Michael Saks and Avi Wigderson. Probabilistic Boolean decision trees and the complexity of evaluating game trees. In Proc. 27th IEEE FOCS, pages 29-38, 1986.
[Sze04] Mario Szegedy. Quantum speed-up of Markov chain based algorithms. In Proc. $45 t h$ IEEE FOCS, pages 32-41, 2004.


[^0]:    *School of Computer Science and Institute for Quantum Computing, University of Waterloo.

[^1]:    ${ }^{1}$ Actually, $\left[\mathrm{ACR}^{+} 07\right.$, Sec. 7$]$ only shows a bound of $O\left(\sqrt{n} d^{3 / 2}\right)$ queries, but this can be improved to $O(\sqrt{n} d)$ using the bounds on $\sigma_{ \pm}(\varphi)$ below [ACR ${ }^{+} 07$, Def. 1]. The improved analysis of the same algorithm in [Rei09b] tightens this to $O(\sqrt{n d})$, which is also the depth-dependence found by [Amb07].

