

Space-Efficient Quantum Error Reduction without log Factors

Aleksandrs Belovs¹ and Stacey Jeffery²

¹Center for Quantum Computing Science, Faculty of Science and Technology, University of Latvia

²QuSoft, CWI & University of Amsterdam, the Netherlands

Given an algorithm that outputs the correct answer with bounded error, say $1/3$, it is sometimes desirable to reduce this error to some arbitrarily small ε – for example, if one wants to call the algorithm many times as a subroutine. The usual method, for both quantum and randomized algorithms, is a procedure called majority voting, which incurs a multiplicative overhead of $\mathcal{O}(\log \frac{1}{\varepsilon})$ from calling the algorithm this many times.

A recent paper [8] introduced a model of quantum computation called *transducers*, and showed how to reduce the “error” of a transducer arbitrarily with only constant overhead, using a construction analogous to majority voting called *purification*. Even error-free transducers map to bounded-error quantum algorithms, so this does not let you reduce algorithmic error for free, but it does allow bounded-error quantum algorithms to be composed without incurring log factors.

In this paper, we present a new highly simplified construction of a purifier, that can be understood as a weighted walk on a line similar to a random walk interpretation of majority voting. Our purifier has much smaller space and time complexity than the previous one. Indeed, it only uses one additional counter, and only performs two increment and two decrement operations on each iteration. It also has quadratically better dependence on the soundness-completeness gap of the algorithm being purified. We prove that our purifier has optimal query complexity, even down to the constant, which matters when one composes quantum algorithms to super-constant depth.

Purifiers can be seen as a way of turning a “Monte Carlo” quantum algorithm into a “Las Vegas” quantum algorithm – a process for which there is no classical analogue – leading to phenomena like the above-mentioned composition of bounded-error quantum algorithms without log factors. Conceptually, our simplified construction sheds light on this strange quantum phenomenon, and, more practically, could have implications for the time and space complexity of composed quantum algorithms.

1 Introduction

Error reduction is one of the most fundamental algorithmic primitives in both the randomized and quantum settings. It transforms a program that evaluates a function f with bounded error, say $1/3$, into a program that evaluates the same function with much smaller error ε . In both the quantum and randomized cases, the conventional procedure is based on majority voting.

If we want to reduce the error of a complete quantum algorithm that ends with a measurement outputting a classical string, then we can treat it as a randomized process and apply majority voting just as we would in a randomized algorithm: execute the quantum algorithm together with the final measurement a number of times, and output the majority.

Somewhat more interesting is error reduction of a *coherent* quantum subroutine. If the subroutine is executed several times, such as when it is employed as an input oracle for another quantum algorithm, then its error must be driven down in order to keep the total error of the whole algorithm sufficiently small.¹ As quantum subroutines may be called in superposition, the error-reduced subroutine cannot use measurements, but must instead be a coherent subroutine.

Let us formulate the settings more formally. We will assume the binary case for most of the paper. This is without loss of generality, as the general case can be obtained from the binary case using the Bernstein-Vazirani trick as we explain in [Section 7](#). We treat the subroutine whose error has to be reduced as an input oracle O to the error reduction procedure. The most natural choice is to take the following *state-generating oracle*

$$O: |0\rangle_{\mathcal{A}}|0\rangle_{\mathcal{W}} \mapsto |\phi\rangle_{\mathcal{AW}} = \sqrt{1-p}|0\rangle_{\mathcal{A}}|\phi_0\rangle_{\mathcal{W}} + \sqrt{p}|1\rangle_{\mathcal{A}}|\phi_1\rangle_{\mathcal{W}}, \quad (1.1)$$

where the first register \mathcal{A} is the output qubit (answer register), and the second register \mathcal{W} contains some internal state of the subroutine (workspace), and both states ϕ_0 and ϕ_1 are normalised. The error reduction procedure R , given this input oracle, must perform the transformation

$$R(O): |0\rangle_{\mathcal{A}} \mapsto \begin{cases} |0\rangle_{\mathcal{A}}, & \text{if } p \leq \frac{1}{2} - \delta; \\ |1\rangle_{\mathcal{A}}, & \text{if } p \geq \frac{1}{2} + \delta; \end{cases} \quad (1.2)$$

with sufficiently high accuracy. Here $2\delta > 0$ is the gap between the accepting and the rejecting probabilities of the subroutine O . That is, the error of the subroutine is $1/2 - \delta$. The space complexity of the subroutine is the number of qubits used by \mathcal{A} and \mathcal{W} .

In the remaining part of the introduction, we first recall the traditional approach to error reduction, that literally performs the transformation in (1.2). Then, we consider approaches that improve on the traditional approach by avoiding the transformation in (1.2). They are based on idealised models of computation like span programs, dual adversary, and transducers.

1.1 Traditional Approach to Error Reduction

The most well-known approach to quantum error reduction is still majority voting.

Theorem 1.1 (Quantum Majority Voting, folklore). *Let $\varepsilon, \delta > 0$ be real parameters. Quantum majority voting ε -approximately performs the transformation (1.2) for any input oracle of the form (1.1). It executes the input oracle O and its inverse O^* each $\ell = O(\frac{1}{\delta^2} \log \frac{1}{\varepsilon})$ times. Its space complexity is $ls + O(\log \ell)$, where s is the space complexity of O . Besides executions of O and O^* , it uses $O(\ell \log \ell)$ other 1- or 2-qubit operations.*

One difference between the classical and quantum case, is that we cannot naïvely reuse the space used by the subroutine in the quantum case, but must instead make ℓ copies of the whole state $|\phi\rangle_{\mathcal{AW}}$ (in contrast to classical majority voting, which can just keep track of how many 1s have been seen so far). This gives the ls term in the space complexity of [Theorem 1.1](#). Let us give a short proof in order to demonstrate this.

¹If the best upper bound we can prove on an algorithm's error probability is $\frac{1}{2}$ or larger, the algorithm is widely considered to be useless.

Proof. Let r be the required output of the algorithm in (1.2), i.e., $r = 0$ if $p \leq \frac{1}{2} - \delta$, and $r = 1$ if $p \geq \frac{1}{2} + \delta$. We also assume that ℓ is a power of 2. The algorithm works as follows. The input oracle O is executed ℓ times, each time on a new pair of registers $\mathcal{A}_i, \mathcal{W}_i$. This gives a state of the form

$$\sum_{b \in \{0,1\}^\ell} (1-p)^{\frac{\ell-|b|}{2}} p^{\frac{|b|}{2}} |b_1\rangle_{\mathcal{A}_1} \cdots |b_\ell\rangle_{\mathcal{A}_\ell} |\phi_{b_1}\rangle_{\mathcal{W}_1} \cdots |\phi_{b_\ell}\rangle_{\mathcal{W}_\ell},$$

where $|b|$ is the Hamming weight of b . Evaluation of the sum of the bits b_1, \dots, b_ℓ into a new register \mathcal{R} gives the state

$$|\psi\rangle = \sum_{b \in \{0,1\}^\ell} (1-p)^{\frac{\ell-|b|}{2}} p^{\frac{|b|}{2}} |b_1\rangle_{\mathcal{A}_1} \cdots |b_\ell\rangle_{\mathcal{A}_\ell} |\phi_{b_1}\rangle_{\mathcal{W}_1} \cdots |\phi_{b_\ell}\rangle_{\mathcal{W}_\ell} ||b\rangle_{\mathcal{R}}.$$

Let $\mathcal{R}_{\text{high}}$ be the most significant qubit of \mathcal{R} . We rewrite the state $|\psi\rangle$ as

$$|\psi\rangle = |r\rangle_{\mathcal{R}_{\text{high}}} |\psi_r\rangle + |1-r\rangle_{\mathcal{R}_{\text{high}}} |\psi_{1-r}\rangle,$$

where $|\psi_r\rangle$ and $|\psi_{1-r}\rangle$ are non-normalized states on the remaining registers. Now we copy the qubit $\mathcal{R}_{\text{high}}$ into the output qubit \mathcal{A} , which gives the state

$$|r\rangle_{\mathcal{A}} |r\rangle_{\mathcal{R}_{\text{high}}} |\psi_r\rangle + |1-r\rangle_{\mathcal{A}} |1-r\rangle_{\mathcal{R}_{\text{high}}} |\psi_{1-r}\rangle, \quad (1.3)$$

and run the whole preceding computation in reverse.

If, instead of (1.3), we had the state

$$|r\rangle_{\mathcal{A}} |r\rangle_{\mathcal{R}_{\text{high}}} |\psi_r\rangle + |r\rangle_{\mathcal{A}} |1-r\rangle_{\mathcal{R}_{\text{high}}} |\psi_{1-r}\rangle, \quad (1.4)$$

we would get the required state $|r\rangle_{\mathcal{A}}$ at the end exactly. Since the difference between the states in (1.3) and (1.4) is of norm $\sqrt{2}\|\psi_{1-r}\|$, the imprecision of the algorithm is also $\sqrt{2}\|\psi_{1-r}\|$ by unitarity (see Section 2.2).

Let us estimate $\|\psi_{1-r}\|$. We consider the case $r = 0$ for concreteness, the second case being similar. The norm squared of ψ_{1-r} equals the probability of $|b| \geq \ell/2$ when b is the sum of ℓ i.i.d. Bernoulli random variables with expectation p . By Hoeffding's inequality:

$$\|\psi_{1-r}\|^2 = \Pr\left[|b| \geq \ell/2\right] \leq \Pr\left[|b| - \mathbb{E}[|b|] \geq \ell\delta\right] \leq e^{-2\ell\delta^2} \leq \frac{\varepsilon^2}{2}$$

for some $\ell = \mathcal{O}\left(\frac{1}{\delta^2} \log \frac{1}{\varepsilon}\right)$.

The registers \mathcal{A}_i and \mathcal{W}_i use ℓs qubits together, and \mathcal{R} uses additional $\mathcal{O}(\log \ell)$. Time complexity of the algorithm is easy to check. \square

Another approach uses quantum amplitude estimation. For this primitive, it is sufficient to use a less powerful *reflecting oracle*:

$$O_{\text{ref}} = O \text{Ref}_{|0\rangle|0\rangle} O^*, \quad (1.5)$$

where $\text{Ref}_{|0\rangle|0\rangle}$ is the reflection about the state $|0\rangle_{\mathcal{A}}|0\rangle_{\mathcal{W}}$ in $\mathcal{A} \otimes \mathcal{W}$. This oracle reflects about the state $|\phi\rangle$ in $\mathcal{A} \otimes \mathcal{W}$. It is easy to implement O_{ref} using two oracle calls and $s + \mathcal{O}(1)$ additional operations. The latter are usually majorized by the time complexity of implementing O , since a subroutine using s qubits will generally have time complexity at least s . On the other hand, implementing O using O_{ref} is complicated. (See [1, 10] for some lower bounds involving the reflecting oracle.)

Even more, it suffices for O_{ref} to act as $O_{\text{ref}}|0\rangle|0\rangle O^*$ only on the span of the vectors $|0\rangle_{\mathcal{A}}|\phi_0\rangle_{\mathcal{W}}$ and $|1\rangle_{\mathcal{A}}|\phi_1\rangle_{\mathcal{W}}$, and it can be arbitrary on its orthogonal complement. That is, we may assume that O_{ref} satisfies the following weaker conditions:

$$O_{\text{ref}} : \begin{cases} \sqrt{1-p}|0\rangle|\phi_0\rangle + \sqrt{p}|1\rangle|\phi_1\rangle \mapsto \sqrt{1-p}|0\rangle|\phi_0\rangle + \sqrt{p}|1\rangle|\phi_1\rangle \\ \sqrt{p}|0\rangle|\phi_0\rangle - \sqrt{1-p}|1\rangle|\phi_1\rangle \mapsto -\sqrt{p}|0\rangle|\phi_0\rangle + \sqrt{1-p}|1\rangle|\phi_1\rangle. \end{cases} \quad (1.6)$$

Theorem 1.2 (Quantum Amplitude Estimation [13]). *Let $\varepsilon, \delta > 0$ be real parameters. There exists a quantum algorithm that, given a copy of the state $|\phi\rangle$ as in (1.1) and query access to an oracle O_{ref} satisfying (1.6), outputs 0 if $p \leq \frac{1}{2} - \delta$ and outputs 1 if $p \geq \frac{1}{2} + \delta$ except with error probability ε . It makes $\ell = \mathcal{O}(\frac{1}{\delta\varepsilon})$ (controlled) queries to the oracle O_{ref} , has space complexity $s + \mathcal{O}(\log \ell)$, and time complexity $\mathcal{O}(\ell \log \ell)$.*

While amplitude estimation is typically stated in terms of calls to the state-generating oracle O from (1.1), it actually uses the reflecting oracle O_{ref} defined as in (1.5), and it is simple to show that it is in fact sufficient for O_{ref} to satisfy (1.6).

Majority voting has better dependence on ε , and amplitude estimation on δ . It is possible to combine the two procedures to obtain the following result.

Corollary 1.3 (Quantum Error Reduction). *Let $\varepsilon, \delta > 0$ be parameters. There exists a quantum algorithm that ε -approximately performs the transformation (1.2) for any input oracle of the form (1.1). It executes the input oracle O and its inverse O^* both $\mathcal{O}(\frac{1}{\delta} \log \frac{1}{\varepsilon})$ times. Its space complexity is $\mathcal{O}(\log \frac{1}{\varepsilon})(s + \log \frac{1}{\delta})$, where s is the space complexity of O . Besides executions of O and O^* , it uses $\tilde{\mathcal{O}}(\frac{1}{\delta} \log \frac{1}{\varepsilon})$ other 1- or 2-qubit operations.*

Proof. We first construct an error reduction algorithm V with constant error $\varepsilon = 1/10$ using Theorem 1.2. It requires one copy of $|\phi\rangle$ and $\mathcal{O}(1/\delta)$ executions of O_{ref} , which can be performed using $\mathcal{O}(1/\delta)$ invocations of O and O^* due to (1.5). The space complexity of V is $s + \mathcal{O}(\log \frac{1}{\delta})$. Its time complexity is $\tilde{\mathcal{O}}(\frac{1}{\delta})$.

To get the imprecision down to ε , we use V as an oracle in Theorem 1.1 with δ being $\Omega(1)$. This requires $\mathcal{O}(\log \frac{1}{\varepsilon})$ executions of V and V^* , and its space complexity is $\mathcal{O}(\log \frac{1}{\varepsilon})$ times the space complexity of V plus negligible $\mathcal{O}(\log \frac{1}{\varepsilon})$. Its time complexity is dominated by $\mathcal{O}(\log \frac{1}{\varepsilon})$ times the time complexity of V . \square

The number of queries made by this algorithm has optimal dependence on ε by [14] as well as on δ by [33]. However, this algorithm is not optimal in terms of memory usage. It is possible to reuse the memory of the subroutine O as shown by Marriott and Watrous in the context of QMA [31] – see also [32] – but this construction still uses additional qubits. Alternatively, one can use Quantum Signal Processing (QSP) [30] for similar results [20]. For instance, we can get the following result, which, to the best of our knowledge, has not been previously stated explicitly.

Theorem 1.4 (QSP-based Error Reduction). *Let $\varepsilon, \delta > 0$ be real parameters. There exists a quantum algorithm that, given query access to any oracle O_{ref} satisfying (1.6), ε -approximately performs the following transformation*

$$|\phi\rangle \mapsto \begin{cases} |\phi\rangle, & \text{if } p \leq \frac{1}{2} - \delta; \\ -|\phi\rangle, & \text{if } p \geq \frac{1}{2} + \delta; \end{cases} \quad (1.7)$$

for any normalised $|\phi\rangle \in \text{span}\{|0\rangle|\phi_0\rangle, |1\rangle|\phi_1\rangle\}$ (including, but not limited to, the specific vector $|\phi\rangle$ from (1.1)). The algorithm makes $\ell = \mathcal{O}(\frac{1}{\delta} \log \frac{1}{\varepsilon})$ queries to the oracle O_{ref} , and uses $\mathcal{O}(\ell)$ other 1- or 2-qubit gates. The algorithm uses no additional qubits compared to O_{ref} .

We prove [Theorem 1.4](#) in [Section 6](#). The proof is similar to a construction in [[20](#), [Section 3](#)], used there for QMA amplification. See [Section 6](#) for a more detailed comparison.

1.2 Span Programs, Dual Adversary, Transducers, and Error Reduction

Given the lower bounds from [[14](#)] and [[33](#)], the algorithm in [Theorem 1.4](#) seems to be the best possible technique for error reduction. Also, often $\delta = \Omega(1)$, so the query complexity is only logarithmic, which might seem negligible. However, imagine a bounded-error subroutine that executes itself recursively many times. In this case, error reduction for each level of recursion gives a multiplicative logarithmic factor, resulting in something like $\mathcal{O}(\log \frac{1}{\varepsilon})^d$, where d is the depth of recursion. If d is super-constant, this might be quite significant.

There exist several approaches to deal with error reduction by avoiding the transformation in [\(1.2\)](#) altogether. They all use similar ideas, and we first explore the one based on span programs [[37](#)], which was directly inspired by quantum algorithms for the iterated NAND tree [[19](#), [16](#), [3](#)]. The span program model is an idealised computational model (meaning that it does not correspond to any actual computational device) for evaluating Boolean functions, with the following three key properties:

1. Any bounded-error quantum query subroutine evaluating a Boolean function can be converted into an exact span program whose complexity (called *witness size*) is the same as the query complexity of the subroutine up to a constant factor.²
2. Span programs can be composed. More precisely, a span program for a function f with witness size W_f can be composed with a span program for a function g with witness size W_g resulting in a span program for the composed function $f \circ g$ with witness size $W_f W_g$.
3. Any span program can be converted into a quantum query algorithm with the same (up to a constant factor) complexity, introducing some small error.

Thus, having two bounded-error quantum query algorithms evaluating functions f and g in Q_f and Q_g queries, respectively, we can first lift them to the idealised (error-free) world of span programs ([Point 1](#)), combine them there ([Point 2](#)), and then bring them back ([Point 3](#)) resulting in a bounded-error quantum query algorithm for the function $f \circ g$ with query complexity $\mathcal{O}(Q_f Q_g)$. This improves on $\mathcal{O}(Q_f Q_g \log Q_g)$, which we would get by composing the programs for f and g directly using [Theorem 1.1](#) or [1.4](#) to reduce the error of the latter. Of course, this approach also works in more complicated settings, like the recursive procedure mentioned in the first paragraph of this section.

[Point 1](#) above holds, for one reason, because span programs just do not have a notion of error.³ The transformation from the quantum query algorithm to the span program in [Point 1](#) is due to [[36](#)] and it is indirect. It goes by proving that span programs are dual to an important quantum lower bound technique: the adversary bound [[22](#)]. Because of the latter, generalisations of span programs beyond evaluation of Boolean functions are generally called *dual adversary bounds*.

²This constant factor is unavoidable here, and it depends on the permitted error probability of the subroutine. We revisit this question in [Section 5](#).

³We are talking about the original version of span programs from [[37](#)] here. There exists a version of span programs capable of approximate evaluation of functions [[23](#)], but this is just for convenience. We know by [[36](#)] that there is always an *exact* span program with optimal witness complexity.

Span programs allow composition without log factors, but they only work for Boolean functions. Further research focused on generalising this to other settings like evaluation of arbitrary functions [28], state generation [4], and state conversion [28]. In [6], a version of the dual adversary for state conversion was constructed that allowed the input oracle to be an arbitrary unitary, not just encoding some input string like in [28]. This made it possible to construct idealised counterparts to the whole spectrum of quantum unitary transformations, and a *direct* reduction from the quantum query algorithm to the dual adversary was shown (building on [28]).

This generality caused some problems with [Point 1](#) above (with span programs replaced by the dual adversary). If one takes a subroutine for evaluating a function, and converts it into the dual adversary for state conversion, the resulting dual adversary will have precisely the same action, and, consequently, precisely the same error! If the algorithm had bounded error, the adversary would have the same error. The remedy proposed in [6] was a *purifier*. It is a dual adversary that has input oracle of the form (1.1), and which performs the transformation (1.2) exactly in constant complexity (assuming $\delta = \Omega(1)$). Of course, the exactness only holds in the idealised world. If one were to directly implement the purifier as a real quantum algorithm (like in [Point 3](#)), this would bring the bounded error back. Composing it with the dual adversary obtained from the bounded-error algorithm above, yields a dual adversary for *exact* evaluation of the function with only $\mathcal{O}(1)$ multiplicative increase in complexity. The purifier can be used in other settings as well. For instance, it can be applied to the input of the algorithm, covering the case when the input oracle is noisy.

The span programs and the dual adversaries described above only consider *query* complexity of quantum algorithms. In [8], the model of *transducers* was introduced. Transducers are a relaxation of both the usual model of quantum algorithms and the dual adversary in the form of [6]. We will describe them in detail in [Section 2](#). For now, let us note that a transducer is a unitary S in a larger space $\mathcal{H} \oplus \mathcal{L}$ that yields a certain unitary $S_{\mathcal{H}}^{\xi}$ on the subspace \mathcal{H} , called the *transduction action* of S . For $\xi, \tau \in \mathcal{H}$, we write $\xi \xrightarrow{S} \tau$ and say that S *transduces* ξ into τ if $S_{\mathcal{H}}^{\xi} \xi = \tau$. The transduction action of the transducer is its idealised action (in the same sense as above), and it can be approximately implemented by executing S several times in a row (see [Theorem 2.6](#) below). More precisely, $\mathcal{O}(W)$ iterations suffice, where W is the so-called *transduction complexity*.

As an example, if S is the walk operator, representing one step of an electric quantum walk [5] on a graph G , and s is the initial vertex of the walk, then we have exact transductions $|s\rangle \xrightarrow{S} |s\rangle$ or $|s\rangle \xrightarrow{S} -|s\rangle$, depending on whether the graph contains marked vertices or not. For details, see [8, Section 6] or [7]. Here S is the easy-to-implement unitary, \mathcal{H} is spanned by $|s\rangle$, and $S_{\mathcal{H}}^{\xi}$ is the action we are interested in.

Since transducers are unitary transformations themselves, they naturally capture time and space complexity as well, not only query complexity.⁴ The *iteration time* complexity $T(S)$ is just the number of gates necessary to implement S (as a usual unitary), and its *space complexity* is the required number of qubits. The query complexity $L(S)$ is defined directly and it is based on the notion of quantum Las Vegas query complexity from [11]. We define the latter formally in [Section 2.1](#), but for now it suffices to note that it can depend on the input oracle and the initial state, and it is always smaller than the number

⁴Span programs also capture space rather naturally [24], but they only capture time complexity in the sense that it is possible to convert a quantum algorithm into an *approximate* span program, and then back to an algorithm, in a way that more or less preserves time complexity – there is no clean way to reason about the time complexity within the span program itself [17].

of queries made by the algorithm.

Again, we can transform any quantum algorithm into a transducer (corresponding to [Point 1](#) above) preserving Las Vegas query complexity, and we can implement the transduction action of any transducer introducing some error (corresponding to [Point 3](#) above). It is also possible to compose transducers in various ways (corresponding to [Point 2](#) above). See [Table 1.1](#) for an informal comparison between complexity measures of a transducer S and an algorithm A performing the same transformation.

Table 1.1

Transducer S	Algorithm A
$\xi \xrightarrow{S} \tau$	$A \xi\rangle \approx \tau\rangle$
$T(S)W(S)$	\leftrightarrow time complexity of A
$L(S)$	\leftrightarrow query complexity of A
space complexity of S	\leftrightarrow space complexity of A

Transducers have some notions of complexity that correspond to time, query and space complexity, and we can map back and forth between bounded-error quantum algorithms and transducers in a way that preserves these complexities. More than that: we can compose transducers to get much better complexity than we would obtain by directly composing the algorithms, which makes them a useful model.

Some remarks related to the exactness condition in [Point 1](#) above are in order. On one hand, as mentioned above, transducers automatically give exact transduction action for any algorithm based on electric quantum walks. This includes such important algorithms as Grover's search [\[21\]](#) and Ambainis' algorithm for element distinctness [\[2\]](#). On the other hand, given an arbitrary quantum algorithm, what we can build is a transducer whose transduction action is *precisely* equal to the action of this algorithm. If the algorithm has bounded error, the transducer inherits it. This is completely analogous to the situation with the dual adversary mentioned above. Ref. [\[8\]](#) applied a similar solution as for the dual adversary and constructed an explicit implementation of the purifier from [\[6\]](#) as a transducer.

1.3 Main Results

In [\[8\]](#), a purifier was constructed as a transducer with query and transduction complexity $\mathcal{O}(1/\delta^2)$. Note that in contrast to majority voting, this is independent of the target error ε , and therefore already allows composition of bounded-error quantum algorithms without log factor overhead. However, the iteration time and space complexity of this purifier were comparable to that of majority voting from [Theorem 1.1](#).

The main contribution of this paper is a better and simpler construction of a purifier. It turns out that the most convenient way to cast our results is in a form similar to the one used by the QSP-based error reduction from [Theorem 1.4](#). The most natural way to state our main result is to use infinite-dimensional space (see [Theorem 3.7](#) for the formal statement of the next theorem):

Theorem 1.5 (Main Result, Infinite-Dimensional Version). *There exists a transducer (purifier) that uses infinite-dimensional space and that, given query access to any input*

oracle O_{ref} satisfying (1.6), performs the following transduction exactly:

$$|\phi\rangle \rightsquigarrow \begin{cases} |\phi\rangle, & \text{if } p < \frac{1}{2}; \\ -|\phi\rangle, & \text{if } p > \frac{1}{2}; \end{cases} \quad (1.8)$$

for any normalised $|\phi\rangle \in \text{span}\{|0\rangle|\phi_0\rangle, |1\rangle|\phi_1\rangle\}$ (including the original vector $|\phi\rangle$ of (1.1)). The query complexity of the transducer is $1/(2\delta)$ and the transduction complexity is $\mathcal{O}(1/\delta)$, where $\delta = |\frac{1}{2} - p|$.

Of course, it is not possible to implement this construction directly. We will explain in a moment how to solve this issue, but for now let us note that our new purifier is extremely simple: besides executions of the input oracle, it only uses two increment and two decrement operations! Like the QSP-based error reduction procedure from Theorem 1.4, it uses the reflecting oracle O_{ref} , and a copy of the state $|\phi\rangle$. Contrary to Theorem 1.4, the purifier need not know δ in advance. In other words, the same transducer works for all values of δ , and the query complexity is input-dependent.

In Section 5, we prove that this transducer is optimal: every transducer with transduction action (1.8) must have query complexity at least $1/(2\delta)$ even if δ is known in advance and there are only two admissible input oracles. Note that this bound is tight *including* the constant factors! The latter are important in this context as they matter for recursions with super-constant depth. On the other hand, the constant factors in the transduction and time complexities are not so important as they do not multiply during composition.

The following corollary (formal version in Corollary 3.8) states the same result in the spirit of Theorem 1.1 with the state-generating input oracle O and the initial state $|0\rangle$.

Corollary 1.6. *There exists a transducer that uses infinite space and, given query access to any input oracle $O: \mathcal{A} \otimes \mathcal{W} \rightarrow \mathcal{A} \otimes \mathcal{W}$ as in (1.1), performs the following transduction exactly:*

$$|0\rangle \rightsquigarrow \begin{cases} |0\rangle, & \text{if } p < \frac{1}{2}; \\ |1\rangle, & \text{if } p > \frac{1}{2}. \end{cases}$$

The query complexity is $1 + 1/(2\delta)$, and the transduction complexity is $\mathcal{O}(1/\delta)$, where $\delta = |\frac{1}{2} - p|$.

The transducer from Theorem 1.5 is exact, but it uses infinite space. Luckily, allowing small perturbations, it is not hard to restrict it to finite (and small!) space. The following result states our main result in the form close to Theorem 1.4 (see Theorem 4.4 for the formal statement):

Theorem 1.7 (Main Result, Informal). *Let $\varepsilon, \delta > 0$ be real parameters. There exists a transducer (purifier) that, given query access to any input oracle O_{ref} satisfying (1.6), performs the following transduction with perturbation ε :*

$$|\phi\rangle \rightsquigarrow \begin{cases} |\phi\rangle, & \text{if } p \leq \frac{1}{2} - \delta; \\ -|\phi\rangle, & \text{if } p \geq \frac{1}{2} + \delta; \end{cases}$$

for any normalised $|\phi\rangle \in \text{span}\{|0\rangle|\phi_0\rangle, |1\rangle|\phi_1\rangle\}$ (including, but not limited to, the specific vector $|\phi\rangle$ from (1.1)). The transducer has query complexity at most $1/(2\delta)$ and transduction complexity $\mathcal{O}(1/\delta)$. It uses $\mathcal{O}(\ell)$ additional qubits and its iteration time is $\mathcal{O}(\ell)$, where $\ell = \log \frac{1}{\delta} + \log \log \frac{1}{\varepsilon}$.

Table 1.2

Complexity	Purifier from [6]	Purifier from [8]	Purifier of Theorem 1.7 (new)
Space	-	$\mathcal{O}(s \frac{1}{\delta^2} \log \frac{1}{\varepsilon})$	$s + \mathcal{O}(\log \frac{1}{\delta} + \log \log \frac{1}{\varepsilon})$
Iteration time	-	$\mathcal{O}(s \frac{1}{\delta^2} \log \frac{1}{\varepsilon})$	$\mathcal{O}(\log \frac{1}{\delta} + \log \log \frac{1}{\varepsilon})$
Query	$\mathcal{O}(\frac{1}{\delta^2})$	$\mathcal{O}(\frac{1}{\delta^2})$	$\frac{1}{2\delta}$
Transduction	-	$\mathcal{O}(\frac{1}{\delta^2})$	$\mathcal{O}(\frac{1}{\delta})$

A side-by-side comparison of our purifier in [Theorem 1.7](#) with the ones from [6] and [8]. The one from [6] is in the form of dual adversary, and only query complexity applies to it. Ref. [6] does not explicitly evaluate the query complexity, only mentioning that it is $\mathcal{O}(1)$ for all $\delta = \Omega(1)$, but it is not hard to work out the required estimate. Iteration time for both transducers in the last two columns is with respect to the circuit model.

Here, the perturbation is a technical term similar to the ε -approximation in [Theorem 1.1](#) and [1.4](#). Also, iteration time here is with respect to the circuit model, that means that time is defined as the number of 1- and 2-qubit operations. Indeed, as we mentioned previously, we only have to make two increment and two decrement operations on each iteration. [Table 1.2](#) provides a side-by-side comparison of the new purifier and the previous ones from [6] and [8]. The purifiers from [6] and [8] use the state-generating input oracle, but this does not affect the asymptotics in the corresponding columns. The inverse quadratic dependence on δ in [6] and [8] was not such a crucial issue, as the purifier could be composed with amplitude estimation like in [Corollary 1.3](#), but that would further increase space and iteration time complexity.

There is another subtle difference between the purifiers from [8] and [Corollary 1.6](#). The purifier from [8] only executed the input oracle on the initial state $|0\rangle$. The purifier from [Corollary 1.6](#) achieves quadratic improvement by executing the input oracle on initial states different from $|0\rangle$. This is not unlike the distinction between majority voting of [Theorem 1.1](#) and amplitude estimation from [Theorem 1.2](#), the latter also obtaining quadratic improvement by executing the input oracle on non- $|0\rangle$ initial states. This is not an issue if the input oracle O is implemented by a usual quantum subroutine, whose complexity is the same from every initial state. On the other hand, if O is a transducer itself, this might cause some complications as it is not sufficient to estimate its complexity only on the initial state $|0\rangle$.

We can also compare the new purifier in [Theorem 1.7](#) with the QSP-based error reduction procedure of [Theorem 1.4](#). The main advantage of the purifier is that its query complexity does not depend on ε . For a fair comparison, it makes sense to assume that the QSP-based error reduction has $\mathcal{O}(\frac{1}{\delta} \log \frac{1}{\varepsilon})$ iterations, each iteration taking $\mathcal{O}(1)$ time, and the number of iterations of the purifier is given by its transduction complexity (see [Theorem 2.6](#)). In this sense, the query complexity and the number of iterations in both approaches has the same inverse linear dependence on δ . The space and iteration time complexities of the QSP-based error reduction is clearly impossible to beat, but these complexities of the new purifier are so small that they are effectively negligible.

The main drawback of the purifier is that in order to get it working, the program and the subroutine have to be converted into transducers and composed similarly to [Point 2](#)

on Page 5. The QSP-based algorithm is advantageous if the overhead of this conversion kills the benefits of smaller query complexity. See also discussion in Section 1.4.

Another minor drawback of transducers is that implementation of their transduction action (like in Point 3 on Page 5, see also Theorem 2.6) has bad inversely quadratic dependence on the required precision ε . Therefore, a natural solution is to implement the transduction action with constant precision, resulting in bounded error, and then either measure the output and use classical majority voting, or, if this is not possible, use the QSP-based error reduction to reduce the error to tolerable value. The benefit is that this error reduction is needed only once for the whole algorithm, and not for each individual subroutine.

1.4 Consequences and Open Problems

One set of consequences of our results are similar to the ones obtained in [8]. We will mention one simple application, Theorem 1.8 below, which is proven in [8] using the old purifier construction. The proofs go through in exactly the same way, only the transducer gets replaced, so we do not repeat them here.

Theorem 1.8 ([8, Theorems 3.15 and 3.16]). *Let A and B be quantum algorithms in the circuit model that evaluate the functions f and g , respectively, with bounded error. Let L be the query complexity of A , and $T(A)$ and $T(B)$ be the time complexities of A and B . Then,*

- (a) *there exists a quantum algorithm in the circuit model that evaluates the composed function $f \circ g$ with bounded error in time complexity*

$$\mathcal{O}(L)(T(A) + T(B)). \tag{1.9}$$

- (b) *There exists a quantum algorithm using quantum random access memory that evaluates the composed function $f \circ g$ with bounded error in time complexity*

$$\mathcal{O}(T(A) + L \cdot T(B)) \tag{1.10}$$

assuming random access and elementary arithmetic operations on memory addresses (word RAM operations) can be done in time $\mathcal{O}(1)$.

Not only does our new purifier improve the space complexity of the composed quantum algorithm referred to in Theorem 1.8, our improved dependence on δ gives better constant-factor overhead, which could matter a lot in algorithms that are obtained from super-constant-depth composition (using extensions of Theorem 1.8 for larger recursion depth). Also, iteration time complexity of our new purifier is so small that it is dominated by other terms in (1.9) and (1.10).

While the expression in (1.10) seems very natural, the one in (1.9) is not so, which is a consequence of the way iteration time behaves under composition of transducers. Eq. (1.9) is optimal if $T(A) \leq T(B)$, though.

Open Problem A. *Is it possible to improve (1.9) to $\mathcal{O}(T(A) + L \cdot T(B))$ in the circuit model?*

Another application of our main result, Theorem 1.5, is a proof that $\text{QMA} = \text{QMA}_1$ assuming that the one-sided-error variant of QMA is allowed to use infinite space [26]. More

generally, it seems quite useful to have an effectively exact representation of bounded-error quantum algorithms, and we suspect this has implications beyond composition.

Let us mention another open problem, of a more technical nature. We demonstrated that the query complexity of the transducer with the reflecting oracle in [Theorem 1.5](#) is the best possible, including the constant factor. For the state-generating variant [Corollary 1.6](#), the picture is not so clear. Since reflection can be implemented using state generation, the former is more powerful than the latter. On the other hand, the latter is more relevant from the practical point of view. It is clear that $\Omega(1/\delta)$ query complexity is necessary, otherwise it would break the lower bound in [\[33\]](#) mentioned above, but the precise value is not clear.

Open Problem B. *Find out the exact value of the query complexity required to solve the problem of [Corollary 1.6](#). The same question can be also asked if we require the purifier to query the state-generating input oracle only on the initial state $|0\rangle$.*


For a high-level didactic overview of some of the concepts of this paper, as well as previous related work, we refer the reader to [\[25\]](#).

1.5 Intuition

In this section, we give an informal overview of the construction behind [Theorem 1.5](#) and its connection to the usual classical majority voting (see also [\[25\]](#)). We construct our purifier as an electric quantum walk on a line. Indeed, as mentioned in [Section 1.2](#), every electric quantum walk gives rise to a transducer whose transduction action is exact, which we take advantage of here.

Before describing our quantum walk, let us revisit classical majority voting. Assume we are given a biased coin that outputs -1 with probability $1-p$, and $+1$ with probability p . We want to decide whether $p \leq \frac{1}{2} - \delta$ (negative case), or $p \geq \frac{1}{2} + \delta$ (positive case). We toss the coin a total of K times. If the sum of all the outcomes is positive, we accept; otherwise, we reject. By Hoeffding's inequality, the error probability is exponentially small in K (similar to [Theorem 1.1](#)).

Figure 1.3



Majority voting can be seen as a random walk on the infinite line.

This can be seen as a random walk on the infinite line (see [Figure 1.3](#)), where we start in the vertex 0 , and, on each step, move to the left with probability $1-p$ and to the right with probability p . The current vertex stores the total sum of all the outcomes. At the end, we accept if we are to the right of the vertex 0 . Using the standard convention that the probability of following an edge is proportional to its weight, this random walk is captured by assigning the weight


$$w_j = \left(\frac{p}{1-p}\right)^j \tag{1.11}$$

to the edge between the vertices j and $j + 1$. Indeed, the probability of going right with this assignment of weights is precisely $p/(1 - p)$ times the probability of going left.

Of course, this random walk still has error, albeit exponentially small. Informally, we get exactness by running this walk for *infinitely long time*. Then, if $p < 1/2$, we will be at $-\infty$; and if $p > 1/2$, we will be at $+\infty$; both with probability 1.

For simplicity of implementation, we cut the line in half, only considering the vertices with non-negative labels. They form an infinite ray (see Figure 1.4). In this graph, the informal exactness statement, made above for the infinite line, can be made precise using *transience* in the sense of Pólya [34] (see also [18, Chapter 2]). If $p < 1/2$, which is equivalent to $p/(1 - p) < 1$, this walk is *recurrent*, meaning that it returns to the vertex 0 infinitely many times with probability 1. If $p > 1/2$, which corresponds to $p/(1 - p) > 1$, this walk is *transient*, meaning that it will escape to infinity with probability 1. Our quantum walk can be seen as deciding which is the case.

Figure 1.4



The quantum walk on the infinite ray used by our purifier. The number above the edge specifies its weight as in (1.11), where we use the same parameter $\gamma = \sqrt{p/(1 - p)}$ we will be using in Section 3 and beyond. Below each edge, the corresponding state of the quantum walk is written.

Let us briefly describe the construction of the quantum walk, more detail being postponed till Section 3.1. The quantum walk happens on edges. The state $|j\rangle$ corresponds to the edge between the vertices j and $j + 1$. The quantum walk is composed of *local reflections*. Each vertex $j > 0$ corresponds to the local reflection about the state

$$\sqrt{w_{j-1}}|j - 1\rangle + \sqrt{w_j}|j\rangle,$$

which is proportional to

$$\sqrt{1 - p}|j - 1\rangle + \sqrt{p}|j\rangle. \quad (1.12)$$

This state is similar to the state ϕ in (1.1), except it does not have the “garbage” states ϕ_0 and ϕ_1 . For simplicity, we will assume that we can perform a reflection about the state (1.12). This is not just to gain intuition: we will make the same assumption in Section 3.1, where we analyse the same quantum walk with the same simplified oracle formally. The general case of Theorem 1.5 turns out to follow quite easily from this special case, see Section 3.2.

Returning to quantum walk, one iteration thereof is given by R_2R_1 , where R_1 is the composition of all local reflections about the odd vertices, and R_2 about the even vertices. To upper bound the number of times the algorithm should call R_2R_1 – that is, the *hitting time* – while ensuring that the algorithm distinguishes between the transient (positive) and recurrent (negative) cases, we need to upper bound the corresponding complexity for each case (for details, see [5]). We will assume here for simplicity that $\delta = \Omega(1)$. The complexity of the quantum walk in the negative case $p \leq \frac{1}{2} - \delta$ is given by the total weight

W of the graph. It equals

$$W = \sum_{j=0}^{\infty} \gamma^{2j} = \mathcal{O}(1)$$

since $\gamma = 1 - \Omega(1)$ in this case.

In the positive case $p \geq \frac{1}{2} + \delta$, the complexity is given by the effective resistance R between the initial and the marked vertices, where resistance of each edge is given by the inverse of its weight. The marked vertex in this case is infinity, hence, the total resistance is the sum of the resistances of all edges:

$$R = \sum_e \frac{1}{w_e} = \sum_{j=0}^{\infty} \frac{1}{\gamma^{2j}} = \mathcal{O}(1)$$

since $\gamma = 1 + \Omega(1)$ in this case.

Quite remarkably, the number of iterations of the quantum walk is finite despite the graph being infinite and the corresponding random walk requiring infinitely many steps to reach $+\infty$. Thus, this gives an instance of a 1-vs-infinity separation between quantum and classical hitting times. This seems to contradict the popular belief that electric quantum walks cannot give better than quadratic speed-up compared to random walks. This belief is rooted in the observation that the classical hitting time is bounded by $\mathcal{O}(WR)$. However, there is a difference between classical and quantum W s and R s. In the classical case, the product WR is computed for each graph, and then maximised over all graphs. In the quantum case, W is maximised over all negative cases, and R over all positive cases, and only then they are multiplied. This might cause significant problems in the design of quantum walks [9]. But in this paper, this distinction works to our advantage.

1.6 Organisation of the Paper

Let us explain the organisation of the remainder of this paper. In [Section 2](#), we formally define and give the necessary background on transducers. In [Section 3](#), we analyse the quantum walk on the infinite ray from [Figure 1.4](#). At first, in [Section 3.1](#), we assume we have access to an input oracle implementing a reflection about the states in [\(1.12\)](#). In [Section 3.2](#), we assume the input oracle O_{ref} from [\(1.6\)](#) with the “garbage” states ϕ_0 and ϕ_1 . It turns out that the garbage states can be taken care of by performing two quantum walks isomorphic to the simple one from [Section 3.1](#).

In [Section 4](#), we describe how the above purifier can be adopted to work in the finite-dimensional settings. First, in [Section 4.1](#), we consider a time-efficient implementation that introduces some small error. Next, in [Section 4.2](#), we describe a variant of the purifier that remains exact, but only works in the settings of the state conversion problem as defined in [\[6\]](#). In [Section 5](#), we show that the query complexity of the purifier from [Theorem 1.5](#) cannot be improved. In [Section 6](#), we describe error reduction procedure based on quantum signal processing, proving [Theorem 1.4](#). Finally, in [Section 7](#), we describe how to extend the results to non-Boolean functions using the Bernstein-Vazirani algorithm [\[12\]](#).

2 Preliminaries

In this section, we describe the key notions on quantum Las Vegas query complexity from [\[11\]](#) and transducers from [\[8\]](#).

2.1 Quantum Query Algorithms and Las Vegas Query Complexity

We define quantum query algorithms following the standard form of alternating queries and unitaries [15], but we allow the input oracle to be an arbitrary unitary [6, 11].

Definition 2.1. A *quantum query algorithm* is a unitary transformation $A = A(O)$ in some space \mathcal{H} that has the form

$$A(O) = U_Q \tilde{O} U_{Q-1} \tilde{O} U_{Q-2} \tilde{O} \cdots U_2 \tilde{O} U_1 \tilde{O} U_0. \quad (2.1)$$

Here, $O: \mathcal{M} \rightarrow \mathcal{M}$ is the *input oracle*, which can be an arbitrary unitary; U_t are arbitrary unitaries in \mathcal{H} , and \tilde{O} is the *query operator* defined as follows.

The space of the algorithm is decomposed as $\mathcal{H} = \mathcal{H}^\circ \oplus \mathcal{H}^\bullet$, where the input oracle acts only on \mathcal{H}^\bullet . For that, we assume $\mathcal{H}^\bullet = \mathcal{H}^\uparrow \otimes \mathcal{M}$ for some space \mathcal{H}^\uparrow . Then, the query operator is

$$\tilde{O} = I^\circ \oplus (I \otimes O), \quad (2.2)$$

where I° and I are the identities on \mathcal{H}° and \mathcal{H}^\uparrow , respectively.

We call the resulting unitary $A(O)$ in (2.1) the *action* of the algorithm (for the particular input oracle), while we usually reserve the term *algorithm* to the specific product on the right-hand side of the same equation.

Not all query algorithms follow this strict form. For instance, it is not always the case that all queries are equal to \tilde{O} . Sometimes, the controls are different, or the input oracle is applied to different registers. However, all such algorithms can be converted into the form (2.1) using standard techniques.

Note that in Definition 2.1 we only consider queries to O , but not to its inverse O^* , i.e., the algorithm has *unidirectional* access to the input oracle. This does not matter for oracles that are self-inverse, like the reflecting oracle O_{ref} from (1.5). However, for other oracles, like the state-generating oracle O from (1.1), we would like to allow *bidirectional* access. We achieve this by replacing the input oracle O with $O \oplus O^*$. This allows the algorithm to apply either O or O^* based on the value of some control register. In this paper, we allow bidirectional access to the state-generating input oracle.

Now we define several complexity measures related to quantum query algorithms. The usual definition of the query complexity of the algorithm A is the number of queries it makes, i.e., Q in (2.1). In particular, it does not depend on the input oracle O , nor the initial state $\xi \in \mathcal{H}$ of the algorithm.

We now define *Las Vegas query complexity* [11], which introduces dependence on both O and ξ . Assume $A(O)$ is executed on some *initial state* $\xi \in \mathcal{H}$. Let


$$\psi_t = U_{t-1} \tilde{O} U_{t-2} \cdots U_1 \tilde{O} U_0 \xi \in \mathcal{H}$$

be the state of the algorithm right before the t -th query. We decompose $\psi_t = \psi_t^\circ \oplus \psi_t^\bullet$ with $\psi_t^\circ \in \mathcal{H}^\circ$ and $\psi_t^\bullet \in \mathcal{H}^\bullet$. Then, the *Las Vegas query complexity* of the algorithm A on the input oracle O and the initial state ξ is defined as

$$L(A, O, \xi) = \sum_{t=1}^Q \|\psi_t^\bullet\|^2. \quad (2.3)$$

In other words, the Las Vegas query complexity is the sum, over all queries, of the squared norms of the parts of the state to which the query is *actually* applied. In particular, $L(A, O, \xi) \leq Q$ for all normalized ξ .

Figure 2.1



A graphical illustration of a quantum query algorithm from (2.1) with $Q = 3$ queries. The algorithm interleaves input-independent unitaries U_t with queries $\tilde{O} = I^\circ \oplus (I \otimes O)$. The intermediate state ψ_t after U_{t-1} and before the t -th query is decomposed as $\psi_t^\circ \oplus \psi_t^\bullet$, where only the second half is processed by the input oracle. (Everywhere in this paper, arrowed wires in parallel represent systems combined via direct sum, rather than tensor product.)

For intuition, note that the Las Vegas (i.e. expected) query complexity of a randomized classical algorithm is the sum over all steps, where a query might be made, of the probability of making a query at that step. Moreover, this notion of quantum query complexity is meaningful in the sense that we can convert A to a (bounded-error) quantum algorithm that makes $\mathcal{O}(L)$ queries, where L is the maximum of $L(A, O, \xi)$ over all admissible pairs of O and ξ . (This follows from Corollary 2.12 and Theorem 2.10 below.)

A more precise description of the work performed by the input oracle is given by the *total query state* defined as

$$q(A, O, \xi) = \sum_{t=1}^Q |t\rangle |\psi_t^\bullet\rangle = \bigoplus_{t=1}^Q \psi_t^\bullet \in \mathbb{C}^Q \otimes \mathcal{H}^\bullet = \mathbb{C}^Q \otimes \mathcal{H}^\dagger \otimes \mathcal{M}. \quad (2.4)$$

In particular, $L(A, O, \xi) = \|q(A, O, \xi)\|^2$. For fixed A and O , the mapping $\xi \mapsto q(A, O, \xi)$ is linear.

The *time complexity* of the algorithm A is defined as the total number of elementary operations (gates) required to implement all the unitaries U_0, U_1, \dots, U_Q used by the algorithm. This complexity depends on the particular model of quantum computation. In this paper we consider the *circuit model*, where elementary operations are all 1- and 2-qubit operations.

2.2 Perturbations

We say that an algorithm performs the required transformation ε -*approximately* or it has *imprecision* ε if the distance between the actual final state of the algorithm and the desired final state is at most ε . The following folklore lemma is useful in estimating imprecision [8].

Lemma 2.2. *Assume we have a collection of unitaries U_1, \dots, U_m all acting in the same vector space \mathcal{H} . Let ψ'_0, \dots, ψ'_m be a collection of vectors in \mathcal{H} such that*

$$\psi'_t = U_t \psi'_{t-1}$$

for all t . Let ψ_0, \dots, ψ_m be another collection of vectors in \mathcal{H} such that $\psi_0 = \psi'_0$ and

$$\|\psi_t - U_t \psi_{t-1}\| \leq \varepsilon_t$$

for all t . Then,

$$\|\psi_m - \psi'_m\| \leq \sum_{t=1}^m \varepsilon_t.$$

In applications of this lemma, U_t stand for sequential sections of a quantum algorithm. The vectors ψ'_t form the sequence of states the algorithm goes through during its execution. The vectors ψ_t form an idealised sequence, which we use instead of ψ'_t in the analysis. More precisely, the analysis looks like

$$\psi_0 \xrightarrow{U_1} U_1 \psi_1 \overset{\varepsilon_1}{\approx} \psi_2 \xrightarrow{U_2} U_2 \psi_2 \overset{\varepsilon_2}{\approx} \psi_3 \xrightarrow{U_3} \dots \overset{\varepsilon_{m-1}}{\approx} \psi_{m-1} \xrightarrow{U_m} U_m \psi_{m-1} \overset{\varepsilon_m}{\approx} \psi_m. \quad (2.5)$$

Here $\overset{\varepsilon}{\approx}$ stands for a (conceptual) *perturbation* of size ε , where we are allowed to arbitrarily move the state of the algorithm by a distance at most ε .

Lemma 2.2 states that m small perturbations in (2.5) result in one big perturbation

$$\psi_m \overset{\varepsilon}{\approx} U_m \cdots U_1 U_0 \psi_0$$

of size $\varepsilon = \sum_{t=1}^m \varepsilon_t$. If ψ_m is the desired final state, the algorithm has imprecision ε .

We can define a *perturbed algorithm* $A = A(O)$ as an algorithm of the form (2.1) where in addition to U_t it is allowed to use perturbations. The total perturbation of the algorithm A is

$$\eta(A, O, \xi) = A(O)\xi - \bar{A}(O)\xi, \quad (2.6)$$

where \bar{A} is the same algorithm as A but with all the perturbations removed.

The analysis like the one in (2.5) is ubiquitous in quantum algorithms, where, for example, perturbations are used after executions of bounded-error subroutines (with errors reduced using [Theorem 1.1](#)) to replace the actual final state of the subroutine with the final state of the ideal subroutine. Introducing perturbations inside the algorithm is helpful as it allows better control on the total query state and the Las Vegas query complexity. In particular, we can avoid subroutines being called on some gibberish initial states whose Las Vegas complexity is hard to estimate. However, we will not pursue this angle in this paper.

2.3 Transducers

The transducer is a new paradigm for quantum algorithms where a single unitary — often simple to implement in some desired sense — gives rise to some target unitary, which we are interested in. A transducer is very easy to define.

Definition 2.3. A *transducer* is a unitary S acting on a direct sum $\mathcal{H} \oplus \mathcal{L}$ of vector spaces, called its *public* and its *private* space, respectively.

We will usually just write “a transducer S in $\mathcal{H} \oplus \mathcal{L}$ ” understanding that the first one is its public, and the second one is its private space. As one can see from the naming, we have a case of encapsulation here: an outside algorithm only has access to the public space. The private space is *only* modified by executions of S .

Definition 2.3 is somewhat empty as it only specifies the easy-to-implement unitary S . The following theorem yields the corresponding target unitary in \mathcal{H} .

Theorem 2.4 ([8, Theorem 5.1]). *Let S be a transducer in $\mathcal{H} \oplus \mathcal{L}$. For any $\xi \in \mathcal{H}$, there exist unique $\tau = \tau(S, \xi) \in \mathcal{H}$ and $v = v(S, \xi) \in \mathcal{L}$ such that*

$$S: \xi \oplus v \mapsto \tau \oplus v. \quad (2.7)$$

Moreover, for a fixed S , the transformation $\xi \mapsto \tau(S, \xi)$ is unitary.

Definition 2.5 (Transduction). For a transducer S in $\mathcal{H} \oplus \mathcal{L}$, its *transduction action* is the unitary $S_{\mathcal{H}}^{\xi}: \mathcal{H} \rightarrow \mathcal{H}$ defined by the mapping $\xi \mapsto \tau(S, \xi)$ from [Theorem 2.4](#). We say that S *transduces* ξ into τ and write $\xi \xrightarrow{S} \tau$ or $S: \xi \rightsquigarrow \tau$. The vector $v(S, \xi)$ is called the *catalyst* of this transduction. Its norm squared $W(S, \xi) = \|v(S, \xi)\|^2$ is called *transduction complexity*. As in [Section 2.1](#), we call ξ the *initial state*, and we use the name *initial coupling* for $\xi \oplus v$.

In general, the catalyst v is not uniquely defined. This is not an issue since we prove that $\xi \xrightarrow{S} \tau$ using specific v that satisfies [\(2.7\)](#), and we call this specific v the catalyst and use it to define transduction complexity (and later query complexity in [Definition 2.7](#)).

The following result states that we can approximately implement the transduction action $S_{\mathcal{H}}^{\xi}$ using black-box access to S .

Theorem 2.6 (Implementation of Transduction Action [8, Theorem 5.5]). *Fix a positive integer K , and spaces \mathcal{H} and \mathcal{L} . There is a quantum algorithm that, for any transducer S in $\mathcal{H} \oplus \mathcal{L}$ and initial state $\xi \in \mathcal{H}$, produces a state τ' such that*

$$\|\tau' - \tau(S, \xi)\| \leq 2\sqrt{\frac{W(S, \xi)}{K}} \quad (2.8)$$

using K controlled calls to S , and $\mathcal{O}(K)$ other elementary operations. The algorithm uses $\mathcal{O}(\log K)$ additional qubits (besides \mathcal{H} and \mathcal{L}).

Proof sketch. The space of the algorithm is $(\mathbb{C}^K \otimes \mathcal{H}) \oplus \mathcal{L}$. The algorithm starts in the initial state $|\xi\rangle$. As the first step, we attach the uniform superposition over K elements, resulting in the state

$$\left(\frac{1}{\sqrt{K}} \sum_{i=1}^K |i\rangle \right) \otimes |\xi\rangle = \bigoplus_{i=1}^K \frac{\xi}{\sqrt{K}} \in \mathbb{C}^K \otimes \mathcal{H}. \quad (2.9)$$

We perform a conceptual perturbation in the sense of [Section 2.2](#) and assume that we are in the state

$$\left(\bigoplus_{i=1}^K \frac{\xi}{\sqrt{K}} \right) \oplus \frac{v}{\sqrt{K}} \quad (2.10)$$

instead. Now we execute S in total K times, each time to map $\frac{\xi}{\sqrt{K}} \oplus \frac{v}{\sqrt{K}}$ into $\frac{\tau}{\sqrt{K}} \oplus \frac{v}{\sqrt{K}}$. On each of these executions, we take a fresh copy of $\frac{\xi}{\sqrt{K}}$ from [\(2.10\)](#) and the same $\frac{v}{\sqrt{K}}$ from the register \mathcal{L} that does not change in the entire process. At the end, we get the state

$$\left(\bigoplus_{i=1}^K \frac{\tau}{\sqrt{K}} \right) \oplus \frac{v}{\sqrt{K}}. \quad (2.11)$$

We perform another conceptual perturbation and assume that we are in the state

$$\bigoplus_{i=1}^K \frac{\tau}{\sqrt{K}} = \left(\frac{1}{\sqrt{K}} \sum_{i=1}^K |i\rangle \right) \otimes |\tau\rangle \quad (2.12)$$

instead. We end the algorithm by detaching the uniform superposition in the first register.

By [Lemma 2.2](#), the imprecision of the algorithm is bounded by the sum of two perturbations — between [\(2.9\)](#) and [\(2.10\)](#), and between [\(2.11\)](#) and [\(2.12\)](#). Each of the perturbations is of size $\left\| \frac{v}{\sqrt{K}} \right\| = \sqrt{\frac{W(S, \xi)}{K}}$. \square

2.4 Complexities of Transducers

A transducer S yields a unitary $S_{\mathcal{H}}$ on the public space \mathcal{H} via [Theorems 2.4](#) and [2.6](#). However, S is also a unitary in $\mathcal{H} \oplus \mathcal{L}$ itself, which can be implemented by a quantum algorithm. We will assume we have some fixed implementation of S by an algorithm, which we will denote by A_S in this section to distinguish from the unitary S per se. (In other parts of the paper, we will usually identify S with its implementation like we do it for A in [Section 2.1](#).)

The transducer $S = S(O)$ can also use an input oracle O . In this case, we may assume that its implementation $A_S = A_S(O)$ is in the form of [Section 2.1](#). To save brackets, we will write $v(S, O, \xi)$ and $\tau(S, O, \xi)$ instead of $v(S(O), \xi)$ and $\tau(S(O), \xi)$, respectively.

Now we can easily define various complexities of a transducer.

Definition 2.7 (Complexities of a transducer). Let S be a transducer. Its *iteration time*⁵, $T(S)$, is the time complexity (number of elementary operations) used by its implementation A_S . Its *space complexity* is the space complexity of A_S .

For fixed input oracle O and initial state ξ , we already defined *transduction complexity* by $W(S, O, \xi) = \|v\|^2$, where $v = v(S, O, \xi)$ is the corresponding catalyst as in [\(2.7\)](#). We also define the *total query state* and the *query complexity* by

$$q(S, O, \xi) = q(A_S, O, \xi \oplus v) \quad \text{and} \quad L(S, O, \xi) = L(A_S, O, \xi \oplus v). \quad (2.13)$$

Note that the state $\xi \oplus v$ above is not normalised, but the definitions in [\(2.3\)](#) and [\(2.4\)](#) still make sense.

Iteration time is precisely where we marry the model-dependent notion of “time complexity” with the nice clean mathematical model of transducers. The notion of time complexity we use for A_S might be with respect to a particular gate set; it might include random access gates (or not); it might be the depth. The time complexity of the algorithm we get in [Theorem 2.6](#) will inherit this model. Note that neither the iteration time nor the space complexity depend on the initial state or the input oracle.

We have the following simple principle that proves the linearity part of [Theorem 2.4](#):

Claim 2.8 (Linearity). *If S transduces ξ_1 into τ_1 with catalyst v_1 and ξ_2 into τ_2 with catalyst v_2 , then S transduces $\alpha\xi_1 + \beta\xi_2$ into $\alpha\tau_1 + \beta\tau_2$ with catalyst $\alpha v_1 + \beta v_2$ for all $\alpha, \beta \in \mathbb{C}$. In particular,*

$$q(S, O, \alpha\xi_1 + \beta\xi_2) = \alpha q(S, O, \xi_1) + \beta q(S, O, \xi_2). \quad (2.14)$$

Also, for every $\xi \in \mathcal{H}$ and $\alpha \in \mathbb{C}$:

$$W(S, O, \alpha\xi) = |\alpha|^2 W(S, O, \xi) \quad \text{and} \quad L(S, O, \alpha\xi) = |\alpha|^2 L(S, O, \xi). \quad (2.15)$$

⁵In [\[8\]](#), this was called time complexity of S , which we now find misleading, as it might suggest that such is the total complexity of the algorithm in [Theorem 2.6](#).

2.5 Canonical Transducers

We mentioned in the previous section, that we may assume that a transducer $S = S(O)$, as an algorithm in $\mathcal{H} \oplus \mathcal{L}$, is in the standard form of (2.1). However, if we naïvely applied the algorithm from Theorem 2.6 to this transducer, this would result in the number of queries $\mathcal{O}(W \cdot Q)$, where W is the transduction complexity and Q is the number of queries made by the transducer.

This can be improved. First, it turns out that transducers admit a much simpler canonical form: it is just one(!) execution of the input oracle on a part of the private space, followed by a unitary independent of the input oracle. Second, there exists a more refined implementation of the transduction action that brings the total number of queries down to the query complexity of the transducer as defined in (2.13). The canonical form of transducers is also important in their composition, which we consider in Section 2.7.

This section is organised as follows. First, we define the canonical form. Then, we state a refined variant of Theorem 2.6 for canonical transducers. Finally, we describe how any transducer can be put into canonical form.

Definition 2.9 (Canonical transducer). A transducer $S = S(O)$ with an input oracle $O: \mathcal{M} \rightarrow \mathcal{M}$ is in *canonical form* if it satisfies the following conditions. First, its private space is decomposed into a direct sum $\mathcal{L} = \mathcal{L}^\circ \oplus \mathcal{L}^\bullet$ with $\mathcal{L}^\bullet = \mathcal{L}^\uparrow \otimes \mathcal{M}$. Second, the transducer itself is in the form

$$S(O) = S^\circ \tilde{O},$$


where $\tilde{O} = I_{\mathcal{H}} \oplus I^\circ \oplus (I \otimes O)$ is the query operator and S° is an arbitrary input-independent *work unitary*. Here I° and I are identities on \mathcal{L}° and \mathcal{L}^\uparrow , respectively.

We call \mathcal{L}° and \mathcal{L}^\bullet the *work* and the *query part* of the private space \mathcal{L} , respectively. Decomposing the catalyst v appropriately: $v = v^\circ \oplus v^\bullet$, the action of S looks like (see also Figure 2.2)

$$S(O) : \xi \oplus v^\circ \oplus v^\bullet \xrightarrow{\tilde{O}} \xi \oplus v^\circ \oplus (I \otimes O)v^\bullet \xrightarrow{S^\circ} \tau \oplus v^\circ \oplus v^\bullet. \quad (2.16)$$

All the vectors (except ξ) in the above equation depend on S , O , and ξ ; i.e, we have $\tau = \tau(S, O, \xi)$, and similarly for v , v° , and v^\bullet .

Figure 2.2



A schematic depiction of a transducer in canonical form. It consists of one application of the input oracle O and an input-independent work unitary S° . The catalyst $v \in \mathcal{L}$ is separated into two parts $v = v^\circ \oplus v^\bullet$ with $v^\circ \in \mathcal{L}^\circ$ and $v^\bullet \in \mathcal{L}^\bullet$. The first one is not processed by the oracle, and the second one is. Note that the input oracle is not applied to the public space.

The total query state and the query complexity take particularly simple form:

$$q(S, O, \xi) = v^\bullet(S, O, \xi) \quad \text{and} \quad L(S, O, \xi) = \|v^\bullet(S, O, \xi)\|^2.$$

The transduction complexity is as before

$$W(S, O, \xi) = \|v(S, O, \xi)\|^2 = \|v^\circ(S, O, \xi)\|^2 + \|v^\bullet(S, O, \xi)\|^2.$$

We have the following refinement of [Theorem 2.6](#), which justifies [Table 1.2](#).

Theorem 2.10 (Refined Implementation of Transduction Action [[8](#), Theorem 3.3]). *Fix spaces \mathcal{H} and $\mathcal{L} = \mathcal{L}^\circ \oplus \mathcal{L}^\bullet$ with $\mathcal{L}^\bullet = \mathcal{L}^\uparrow \otimes \mathcal{M}$, and positive real parameters W , L , and ε . Then there exists a quantum algorithm with query access $O: \mathcal{M} \rightarrow \mathcal{M}$ that, for any canonical transducer $S = S(O)$ in $\mathcal{H} \oplus \mathcal{L}$ and any initial state $\xi \in \mathcal{H}$, transforms the state ξ into τ' such that*

$$\|\tau' - \tau(S, O, \xi)\| \leq \varepsilon$$

assuming that $W(S, O, \xi) \leq W$ and $L(S, O, \xi) \leq L$. The algorithm uses $\mathcal{O}(L(S, O, \xi)/\varepsilon^2)$ queries to O , $K = \mathcal{O}(1 + W(S, O, \xi)/\varepsilon^2)$ controlled calls to S° , $\mathcal{O}(K)$ other elementary operations, and $\mathcal{O}(\log K)$ additional qubits.

Thus, in particular, the time complexity of the algorithm described in [Theorem 2.10](#) (including executions of S° , but excluding executions of O) is $\mathcal{O}(1 + W(S, O, \xi)/\varepsilon^2) \cdot T(S)$, where $T(S)$, is the iteration time of S , which, in agreement to [Definition 2.7](#), is the time complexity of the algorithm implementing S° .

When designing a transducer, treating it as a usual quantum query algorithm from [Section 2.1](#) is very convenient, and we will continue doing so in this paper. It is not hard to transform any such transducer into the canonical form with a slight increase in complexity.

Theorem 2.11 (Conversion to Canonical Form [[8](#), Proposition 10.4]). *Assume $S = S(O)$ is a non-canonical transducer with public space \mathcal{H} making Q queries to the input oracle O . We also assume that S , as a query algorithm, is strictly in the form of [\(2.1\)](#), in particular, all its queries to the input oracle are identical. Then, there exists a canonical transducer $S' = S'(O)$ with the same public space \mathcal{H} and the same input oracle that has exactly the same transduction action as S . Moreover,*

$$q(S', O, \xi) = q(S, O, \xi), \quad L(S', O, \xi) = L(S, O, \xi),$$

and

$$W(S', O, \xi) = W(S, O, \xi) + L(S, O, \xi).$$

The transducer S' uses $\mathcal{O}(\log Q)$ additional qubits and its iteration time is at most $\mathcal{O}(1)$ times the iteration time of S .

Note that any algorithm $A = A(O)$ as in [\(2.1\)](#) can be considered as a transducer $S_A = S_A(O)$ with the same transduction action if we just let the private space \mathcal{L} to be empty and assume $v = 0$ in [\(2.7\)](#). This gives the following corollary of the above theorem:

Corollary 2.12. *Assume A is a quantum query algorithm as in [\(2.1\)](#), working in space \mathcal{H} and making Q queries. Then, there exists a canonical transducer S_A with the public space \mathcal{H} , whose transduction action is identical to the action of A . Moreover, for any $\xi \in \mathcal{H}$ and input oracle O :*

$$q(S_A, O, \xi) = q(A, O, \xi) \quad \text{and} \quad W(S_A, O, \xi) = L(S_A, O, \xi) = L(A, O, \xi).$$

Iteration time complexity of S_A is equal to the time complexity of A up to a constant factor, and S_A uses $\mathcal{O}(\log Q)$ additional qubits.

2.6 Relation to the Adversary Bound

In this section, we describe the relation between transducers and the dual adversary bound for state conversion [28, 6]. Before we describe the connection, we have to define the latter two objects.

Definition 2.13. A *state conversion problem* is specified by two Hilbert spaces \mathcal{M} and \mathcal{H} together with a collection of unitaries O_x in \mathcal{M} and pairs of states $\xi_x \mapsto \tau_x$ in \mathcal{H} , where x ranges over some set X . We say that an algorithm A (respectively, transducer S) *solves* the problem (exactly) if $A(O_x)\xi_x = \tau_x$ (respectively, $\xi_x \xrightarrow{S(O_x)} \tau_x$) for all $x \in X$.

Definition 2.14. The *dual adversary bound* corresponding to the state conversion problem of Definition 2.13 is given by the optimisation problem to minimise $\max_{x \in X} \|v_x\|^2$, where $v_x \in \mathcal{H}^\uparrow \otimes \mathcal{M}$, as x ranges over X , are some vectors satisfying the constraint

$$\langle \xi_x, \xi_y \rangle - \langle \tau_x, \tau_y \rangle = \langle v_x, I \otimes (I_{\mathcal{M}} - O_x^* O_y) v_y \rangle \quad \text{for all } x, y \in X. \quad (2.17)$$

Here \mathcal{H}^\uparrow can be an arbitrary space, and I is the identity on that space.

Note that (2.17) can be rewritten in a possibly friendlier equivalent form

$$\langle \xi_x, \xi_y \rangle - \langle \tau_x, \tau_y \rangle = \langle v_x, v_y \rangle - \langle (I \otimes O_x) v_x, (I \otimes O_y) v_y \rangle \quad \text{for all } x, y \in X. \quad (2.18)$$

The optimal value of the dual adversary problem lower bounds the worst-case query complexity of any algorithm solving the corresponding state conversion problem. This is a corollary of the following result connecting dual adversary to transducers.

Theorem 2.15. Let ξ_x, τ_x, O_x for $x \in X$ define a state conversion problem as in Definition 2.13, and let v_x for $x \in X$ be vectors in $\mathcal{H}^\uparrow \otimes \mathcal{M}$. Consider the following statements:

- (a) There exists a quantum query algorithm $A = A(O)$ such that $A(O_x)\xi_x = \tau_x$ and $v_x = q(A, O_x, \xi_x)$ for all $x \in X$.
- (b) There exists a transducer $S = S(O)$ such that $\xi_x \xrightarrow{S(O_x)} \tau_x$ and $v_x = q(S, O_x, \xi_x)$ for all $x \in X$.
- (c) There exists a canonical transducer satisfying (b).
- (d) The vectors v_x satisfy the constraints (2.18).

Then, the statements (b), (c), and (d) are equivalent. The statement (a) implies all of them.

Proof. The implication (a) \implies (c) is the content of Corollary 2.12, while the implication (b) \implies (c) is the content of Theorem 2.11. Implication (c) \implies (b) is trivial as any canonical transducer is a transducer. It remains to establish equivalence of (c) and (d).

Assume there exists a canonical transducer as in Definition 2.9 performing the required transduction. Then, by (2.16), there exists a unitary S° such that

$$S^\circ : \xi_x \oplus v_x^\circ \oplus (I \otimes O_x) v_x^\bullet \longmapsto \tau_x \oplus v_x^\circ \oplus v_x^\bullet$$

for all $x \in X$, where v_x^\bullet are the corresponding total query states. Existence of such unitary is equivalent to the condition

$$\langle \xi_x, \xi_y \rangle + \langle v_x^\circ, v_y^\circ \rangle + \langle (I \otimes O_x) v_x^\bullet, (I \otimes O_y) v_y^\bullet \rangle = \langle \tau_x, \tau_y \rangle + \langle v_x^\circ, v_y^\circ \rangle + \langle v_x^\bullet, v_y^\bullet \rangle \quad \text{for all } x, y \in X,$$

which implies (2.18) with v_x replaced by v_x^\bullet .

On the other hand, if (2.18) is satisfied, then there exists a unitary S° such that

$$S^\circ: \xi_x \oplus (I \otimes O_x)v_x \longmapsto \tau_x \oplus v_x \quad (2.19)$$

for all $x \in X$. This is an instance of a canonical transducer with \mathcal{L}° being empty, v_x° equal to zero, and where the total query states are given by v_x . \square

This result immediately implies the following corollary.

Corollary 2.16. *The optimal value of the dual adversary problem lower bounds all of the following:*

- (a) *Worst-case query complexity $\max_{x \in X} L(S, O_x, \xi_x)$ of any transducer S such that $\xi_x \xrightarrow{S(O_x)} \tau_x$ for all x .*
- (b) *Worst-case Las Vegas query complexity $\max_{x \in X} L(A, O_x, \xi_x)$ of any algorithm A such that $A(O_x)\xi_x = \tau_x$ for all $x \in X$.*
- (c) *The number of queries made by any algorithm from (b), assuming $\|\xi_x\| = 1$ for all x .*

Lower bounding the optimal value of the dual adversary is the content of the usual (primal) adversary bound. Note though that [Theorem 2.15](#) gives much tighter connection between the dual adversary and the transducers solving the corresponding state conversion problem than just worst-case Las Vegas complexity.

[Corollary 2.16](#) only considers algorithms and transducers solving the state conversion problem *exactly*. In [Section 4.2](#), we briefly explain how to use purifiers in order to lower bound complexity of *bounded-error* algorithms in the case of function evaluation.

2.7 Composition of Transducers

We can compose transducers similar to how we can compose quantum algorithms, and the complexity of these compositions follows some natural rules. In this section, we cover parallel and functional composition. [Ref. \[8\]](#) also considers sequential composition.

Algorithms To get some intuition, we start with composition of query algorithms as in [Section 2.1](#). We first consider parallel composition, which corresponds to performing algorithms in superposition. Let $A_i = A_i(O_i)$ with $O_i: \mathcal{M}_i \rightarrow \mathcal{M}_i$ be quantum algorithms for $i = 1, \dots, m$. Their *parallel composition* $\bigoplus_i A_i$ is a quantum algorithm with input oracle O acting in $\bigoplus_i \mathcal{M}_i$ whose action satisfies

$$\left(\bigoplus_i A_i\right)\left(\bigoplus_i O_i\right) = \bigoplus_i A_i(O_i) \quad (2.20)$$

for all unitaries $O_i: \mathcal{M}_i \rightarrow \mathcal{M}_i$. Its total query state is

$$q\left(\bigoplus_i A_i, \bigoplus_i O_i, \bigoplus_i \xi_i\right) = \bigoplus_i q(A_i, O_i, \xi_i). \quad (2.21)$$

If all A_i are equal to some A , we can replace $\bigoplus_i A_i$ with $I_m \otimes A$, however, we often drop I_m in this context. If all O_i are equal to some O , we can replace $\bigoplus_i O_i$ with just O .

(It is easy to see from (2.2) that O and $I_m \otimes O$ are equally powerful.) Finally, the most important case is when all A_i are equal, and all O_i are equal. In this case, we just write

$$q(A, O, \xi) \text{ instead of } q(I_{\mathcal{E}} \otimes A, O, \xi) \quad (2.22)$$

for every $\xi \in \mathcal{E} \otimes \mathcal{H}$ in arbitrary \mathcal{E} , since the space \mathcal{E} can be deduced from ξ . We use the same convention for $L(A, O, \xi)$.

Now let us move to functional composition, where one algorithm implements the input oracle of another algorithm. Assume A is a quantum algorithm as in (2.1) with the input oracle $O: \mathcal{M} \rightarrow \mathcal{M}$, and let B be an algorithm that acts in \mathcal{M} . Their *functional composition* $A \circ B$ is the algorithm A , where each execution of O is replaced by execution of B . The action of $A \circ B$ is $A(B)$, and it does not have the input oracle.

The above description is simple, but it omits a large number of important details. First, we usually have some global input oracle, which is unrelated to the subroutine B . From now on, we will denote it by $O: \mathcal{M} \rightarrow \mathcal{M}$. Thus, the algorithm A has access to two input oracles: the global input oracle O , and $O': \mathcal{M}' \rightarrow \mathcal{M}'$ that gets replaced by B . The subroutine B now acts in \mathcal{M}' , and it also has access to O , as does the functional composition $A \circ B$.

Modelling A 's oracle access to two input oracles O and O' in the framework of Section 2.1 is easy. We assume it has access to one combined input oracle $O \oplus O'$ acting in $\mathcal{M} \oplus \mathcal{M}'$. We can also measure the query complexity of each of the oracles separately by decomposing the total query state from (2.4) as

$$q(A, O \oplus O', \xi) = q^{(0)}(A, O \oplus O', \xi) \oplus q^{(1)}(A, O \oplus O', \xi),$$

where $q^{(0)}(A, O \oplus O', \xi) \in \mathbb{C}^Q \otimes \mathcal{H}^\uparrow \otimes \mathcal{M}$ and $q^{(1)}(A, O \oplus O', \xi) \in \mathbb{C}^Q \otimes \mathcal{H}^\uparrow \otimes \mathcal{M}'$ denote the parts processed by the oracles O and O' , respectively.

With these conventions, we can write the action of $A \circ B$ as $A(O \oplus B(O))$. It is also not hard to describe the total query state of $A \circ B$ by the following expression [11]:

$$q(A \circ B, O, \xi) = q^{(0)}(A, O \oplus B(O), \xi) \oplus q\left(B, O, q^{(1)}(A, O \oplus B(O), \xi)\right). \quad (2.23)$$

Although this expression is lengthy, it is quite natural. It says that the queries made by $A \circ B$ fall into the queries made by A to O directly (the first term), and the queries made by the subroutine B (the second term). For the latter, we use the convention (2.22), and we assume that B is executed on the entire total query state in one go, instead of dividing it into Q chunks as A does. Finally, $B(O)$ replaces O' since B implements this input oracle.

Transducers Both parallel and functional composition of transducers mimics the case of algorithms very closely. Parallel composition $\bigoplus_i S_i$ of transducers $S_i = S_i(O_i)$ can be implemented as parallel composition of their implementations $\bigoplus A_{S_i}$, where we used notation of Section 2.4. For canonical transducers, it boils down to implementing $\bigoplus S_i^\circ$, where S_i° are the corresponding work unitaries. Similarly to (2.21), we have

$$q\left(\bigoplus_i S_i, \bigoplus_i O_i, \bigoplus_i \xi_i\right) = \bigoplus_i q(S_i, O_i, \xi_i) \quad (2.24)$$

and

$$W\left(\bigoplus_i S_i, \bigoplus_i O_i, \bigoplus_i \xi_i\right) = \sum_i W(S_i, O_i, \xi_i). \quad (2.25)$$

For every transducer $S = S(O)$ and every space \mathcal{E} , we can define transducer $I_{\mathcal{E}} \otimes S$, which corresponds to the special case of parallel composition when all S_i are equal to S , and all O_i are equal to O . As in (2.22), we will write, e.g., $q(S, O, \xi)$ instead of $q(I_{\mathcal{E}} \otimes A, O, \xi)$ for every $\xi \in \mathcal{E} \otimes \mathcal{H}$.

Functional composition is slightly more subtle, but similar. Let S_A and S_B be canonical transducers with public spaces \mathcal{H} , and \mathcal{M}' , respectively. We assume S_A has input oracle $O \oplus O'$ acting in $\mathcal{M} \oplus \mathcal{M}'$, and S_B has input oracle O acting in \mathcal{M} . Let $U_A(O \oplus O')$ and $U_B(O)$ denote the transduction actions of S_A and S_B on \mathcal{H} and \mathcal{M}' , respectively. As for algorithms, we can decompose the total query state of A as

$$q(A, O \oplus O', \xi) = q^{(0)}(A, O \oplus O', \xi) \oplus q^{(1)}(A, O \oplus O', \xi)$$

into the parts processed by O and O' , respectively.

Theorem 2.17. *In the above assumptions, there exists a canonical transducer $S_A \circ S_B$ on the public space \mathcal{H} and with the input oracle $O: \mathcal{M} \rightarrow \mathcal{M}$. Transduction action of $S_A \circ S_B$ equals $U_A(O \oplus U_B(O))$. The total query state of $S_A \circ S_B$ satisfies*

$$q(S_A \circ S_B, O, \xi) = q^{(0)}(S_A, O \oplus U_B(O), \xi) \oplus q\left(S_B, O, q^{(1)}(S_A, O \oplus U_B(O), \xi)\right), \quad (2.26)$$

and the transduction complexity is

$$W(S_A \circ S_B, O, \xi) = W(S_A, O \oplus U_B(O), \xi) + W\left(S_B, O, q^{(1)}(S_A, O \oplus U_B(O), \xi)\right). \quad (2.27)$$

The iteration time of $S_A \circ S_B$ is $T(S_A) + T(S_B)$, and its space complexity is the sum of space complexities of S_A and S_B .

The logic behind formulae (2.26) and (2.27) is the same as behind (2.23).

2.8 Perturbed Transducers

In Section 2.4, we assumed that the algorithm A_S implementing S does its job exactly. Most of the results in this section hold if we allow small perturbations in the implementation of S .

A *perturbed transducer* S is a transducer whose implementation A_S is a perturbed algorithm in the sense of Section 2.2. The requirement for S to transduce ξ into τ is still that S maps $\xi \oplus v$ into $\tau \oplus v$ (with the help of the perturbations). We define the perturbation of the transducer as

$$\eta(S, O, \xi) = \eta(A_S, O, \xi \oplus v),$$

where v is the corresponding catalyst, and the right-hand side is as in (2.6).

All the definitions made for usual transducers can be adopted for perturbed transducers, and the implementations in Theorem 2.6 and 2.10 hold for them as well. For example, the right-hand side of (2.8) in Theorem 2.6 estimates the perturbations used in implementation of the transduction action, and Eq. (2.8) still holds (assuming each execution of S in the theorem is still an execution of the perturbed version of S). The perturbations in S amount to K total perturbations, each equal to $\eta(S, O, \xi)/\sqrt{K}$, which allows to estimate the total perturbation of the resulting algorithm using Lemma 2.2. It is also possible to estimate the perturbation of the composed transducer like in Section 2.7, but will not go into detail as it is outside the scope of this paper. For more detail, see [8].

3 Infinite-Dimensional Construction

In this section, we describe a version of the purifier in infinite-dimensional space. In [Section 3.1](#), we start with a formal analysis of the quantum walk on the infinite ray from [Section 1.5](#), which uses a simplified input oracle. In [Section 3.2](#), we generalize it to the real input oracle O_{ref} from (1.6), thus proving [Theorem 1.5](#). In [Section 3.3](#), we prove [Corollary 1.6](#).

Although the purifiers constructed in this section cannot be implemented directly due to their infinite space requirements, they can still be analysed and they demonstrate all the main ideas in the clearest way. In [Section 4](#), we will modify the purifiers so that they use finite space.

3.1 Special Case

In this section, we perform a formal analysis of the quantum walk on the infinite ray described in [Section 1.5](#). As there, we assume a simplified input oracle $O_p: \mathbb{C}^2 \rightarrow \mathbb{C}^2$ reflecting about the state

$$\varphi_p = \sqrt{1-p}|0\rangle + \sqrt{p}|1\rangle, \quad (3.1)$$

which is similar to the state in (1.1), but without the “garbage” parts ϕ_0 and ϕ_1 . In this section, we use parameter

$$\gamma = \sqrt{\frac{p}{1-p}}. \quad (3.2)$$


The two extreme cases $p = 0$ and $p = 1$, corresponding to $\gamma = 0$ and $\gamma = \infty$, respectively, are also covered by the analysis, where we use the usual convention that $0^{-1} = \infty$ and $\infty^{-1} = 0$.

The purifier transducer $S = S(O_p)$ acts in one register \mathcal{N} which stores a non-negative integer, usually denoted by j in this section. However, to make the increment and the decrement operations in \mathcal{N} unitary, we have to assume that \mathcal{N} can contain an arbitrary integer, both positive or negative. We use \mathcal{A} to denote the least significant qubit of \mathcal{N} , and \mathcal{K} for the remaining qubits of \mathcal{N} . Thus, $|i\rangle_{\mathcal{K}}|b\rangle_{\mathcal{A}} = |2i + b\rangle_{\mathcal{N}}$ for an integer i and $b \in \{0, 1\}$. The public space of S is spanned by $|0\rangle_{\mathcal{N}}$, and the private space by $|j\rangle_{\mathcal{N}}$ with $j > 0$. As mentioned above, the states $|j\rangle_{\mathcal{N}}$ with $j < 0$ are not used. The transducer

$$S = R_2 R_1, \quad (3.3)$$

is a product of two reflections as defined in [Figure 3.1](#).

Figure 3.1



Two reflections comprising the transducer $S = R_2 R_1$. Both reflections act in the register $\mathcal{N} = \mathcal{A} \otimes \mathcal{K}$, where \mathcal{A} stores the least significant qubit of \mathcal{N} . The $+1$ and the -1 operators increment and decrement the register \mathcal{N} , respectively. The control for O_p in R_2 is on \mathcal{K} being non-zero.

As we will shortly see in [Claim 3.1](#), R_1 and R_2 , restricted to the span of $|j\rangle_{\mathcal{N}}$ for $j \geq 0$, are the local reflections around the odd and even vertices in [Figure 1.4](#). This includes the local reflection corresponding to the vertex 0, which reflects about the state $|0\rangle$.

The following vectors will be important in the analysis, where j is a positive integer:

$$\begin{aligned} u_j &= |j-1\rangle + \gamma|j\rangle & \text{and} & & u_j^\perp &= |j-1\rangle - \gamma^{-1}|j\rangle \\ &\propto \sqrt{1-p}|j-1\rangle + \sqrt{p}|j\rangle & & & &\propto \sqrt{p}|j-1\rangle - \sqrt{1-p}|j\rangle. \end{aligned} \quad (3.4)$$

In the extreme case of $\gamma = \infty$, we define $u_j = |j\rangle$ and $u_j^\perp = |j-1\rangle$, which still satisfy the proportionality claims in [\(3.4\)](#). Similarly, if $\gamma = 0$, we let $u_j^\perp = -|j\rangle$, which also satisfies the proportionality claim.

Claim 3.1. *For all odd j , we have $R_1 u_j = u_j$ and $R_1 u_j^\perp = -u_j^\perp$. For all even $j > 0$, we have $R_2 u_j = u_j$ and $R_2 u_j^\perp = -u_j^\perp$. Also, $R_2 |0\rangle = |0\rangle$.*

Proof. By construction, for every non-negative integer i , R_1 acts as identity on $|i\rangle_{\mathcal{K}} \otimes |\varphi_p\rangle_{\mathcal{A}} = \sqrt{1-p}|2i\rangle + \sqrt{p}|2i+1\rangle$, and flips the phase of (multiplies by -1) its orthogonal complement in the span of $|2i\rangle$ and $|2i+1\rangle$. The state u_{2i+1} is proportional to the former, and u_{2i+1}^\perp to the latter. The case of R_2 is proven similarly. \square

We can now prove that $S = R_2 R_1$ is an *exact* transducer with the desired transduction action, and query complexity precisely $1/(2\delta) = |1 - 2p|^{-1}$.

Theorem 3.2. *In the above notation, $S = R_2 R_1$ performs the following transduction:*

$$S(O_p): |0\rangle \rightsquigarrow \begin{cases} |0\rangle, & \text{if } p < \frac{1}{2}; \\ -|0\rangle, & \text{if } p > \frac{1}{2}. \end{cases} \quad (3.5)$$

The query complexity of the transducer is $L(S, O_p, |0\rangle) = 1/(2\delta)$, and the transduction complexity is $W(S, O_p, |0\rangle) = \mathcal{O}(1/\delta)$, where $\delta = |\frac{1}{2} - p|$.

Proof. If $p < \frac{1}{2}$, then $\gamma < 1$, and we define the catalyst $|v\rangle = \sum_{j=1}^{\infty} \gamma^j |j\rangle$. The action of the transducer on the initial coupling is given by

$$\begin{aligned} |0\rangle + |v\rangle &= \sum_{j=0}^{\infty} \gamma^j |j\rangle = \sum_{i=0}^{\infty} \gamma^{2i} (|2i\rangle + \gamma|2i+1\rangle) \\ &\xrightarrow{R_1} \sum_{j=0}^{\infty} \gamma^j |j\rangle = |0\rangle + \sum_{i=0}^{\infty} \gamma^{2i+1} (|2i+1\rangle + \gamma|2i+2\rangle) \\ &\xrightarrow{R_2} \sum_{j=0}^{\infty} \gamma^j |j\rangle = |0\rangle + |v\rangle, \end{aligned} \quad (3.6)$$

where we used [Claim 3.1](#). Therefore, $|0\rangle \rightsquigarrow^{S(O_p)} |0\rangle$. The input oracle is applied to the states in the brackets above, therefore, by [\(2.3\)](#), the Las Vegas query complexity is given by

$$\begin{aligned} \sum_{i=0}^{\infty} \left\| \gamma^{2i} (|2i\rangle + \gamma|2i+1\rangle) \right\|^2 + \sum_{i=0}^{\infty} \left\| \gamma^{2i+1} (|2i+1\rangle + \gamma|2i+2\rangle) \right\|^2 \\ = \sum_{j=0}^{\infty} \gamma^{2j} + \sum_{j=1}^{\infty} \gamma^{2j} = \frac{2}{1-\gamma^2} - 1 = \frac{1}{2\delta}. \end{aligned}$$

Now, if $p > \frac{1}{2}$, then $\gamma > 1$, and we define $|v\rangle = \sum_{j=1}^{\infty} (-\gamma)^{-j} |j\rangle$. The action of the transducer on the initial coupling is given by

$$\begin{aligned}
|0\rangle + |v\rangle &= \sum_{j=0}^{\infty} (-\gamma)^{-j} |j\rangle = \sum_{i=0}^{\infty} \gamma^{-2i} (|2i\rangle - \gamma^{-1} |2i+1\rangle) \\
&\xrightarrow{R_1} - \sum_{j=0}^{\infty} (-\gamma)^{-j} |j\rangle = -|0\rangle + \sum_{i=0}^{\infty} \gamma^{-2i-1} (|2i+1\rangle - \gamma^{-1} |2i+2\rangle) \\
&\xrightarrow{R_2} -|0\rangle - \sum_{i=0}^{\infty} \gamma^{-2i-1} (|2i+1\rangle - \gamma^{-1} |2i+2\rangle) \\
&= -|0\rangle + \sum_{j=1}^{\infty} (-\gamma)^{-j} |j\rangle = -|0\rangle + |v\rangle.
\end{aligned} \tag{3.7}$$

Therefore, $|0\rangle \xrightarrow{S(O_p)} -|0\rangle$. The Las Vegas query complexity is

$$\sum_{j=0}^{\infty} \gamma^{-2j} + \sum_{j=1}^{\infty} \gamma^{-2j} = \frac{2}{1 - \gamma^{-2}} - 1 = \frac{1}{2\delta}.$$

By a similar computation, we can show that $\|v\|^2 \leq \frac{1}{2\delta}$ in both cases, showing that the transduction complexity is at most $\frac{1}{2\delta}$. \square

Remark 3.3. Note that [Theorem 3.2](#) holds for $S = R_1 R_2$ as well. The sequence [\(3.6\)](#) is similar with the same catalyst; in the sequence [\(3.7\)](#), we should replace the catalyst with $-\sum_{j=1}^{\infty} (-\gamma)^{-j} |j\rangle$.

3.2 General Case

In this section, we prove [Theorem 1.5](#). For that, we define the purifier transducer $S' = S'(O_{\text{ref}})$ for the general case when the reflecting oracle O_{ref} in $\mathcal{A} \otimes \mathcal{W}$ satisfies [\(1.6\)](#). We first describe the transducer, and then explain its relation to the purifier from [Section 3.1](#).

The transducer uses the register \mathcal{J} for storing an index, which can be an arbitrary non-negative integer. Again, in order to make the increment and decrement unitary, we assume that \mathcal{J} can store an arbitrary integer, but we will not use the negative values. The transducer also uses the registers \mathcal{A} and \mathcal{W} acted on by the input oracle. Note that this time, contrary to [Section 3.1](#), \mathcal{A} is *not* a part of the register \mathcal{J} , but a separate register. When describing the states, the registers are listed in this order. The public space of the transducer is $|0\rangle_{\mathcal{J}} \otimes \mathcal{A} \otimes \mathcal{W}$, which we identify with $\mathcal{A} \otimes \mathcal{W}$. The purifier transducer is given by

$$S' = R'_2 R'_1 \tag{3.8}$$

for reflections R'_1 and R'_2 defined as in [Figure 3.2](#). Each of R'_1 and R'_2 makes precisely one query to O_{ref} , thus, S' makes two queries.


The implementation [\(3.8\)](#) of S' is in the form of [\(2.1\)](#), which can be achieved by breaking the controlled- $(-O_{\text{ref}})$ operation in R'_2 into the controlled phase and the controlled- O_{ref} operations.

The following claim is similar to [Claim 3.1](#).

Claim 3.4. *The operator R'_1 acts as identity on the states*

$$|0\rangle |1\rangle |\phi_1\rangle \quad \text{and} \quad |j-1\rangle |0\rangle |\phi_0\rangle + \gamma |j\rangle |1\rangle |\phi_1\rangle \tag{3.9}$$

Figure 3.2



Two reflections that give the transducer $S' = R'_2 R'_1$. Closed controls (\bullet), including the one from the register \mathcal{J} , are controlled on the wire being non-zero, whereas open controls (\circ) are controlled on the wire being 0. The gates $+1$ and -1 denote increment and decrement in \mathcal{J} , respectively.

for all positive integers j . The following states get reflected for all positive integers j :

$$|j-1\rangle|0\rangle|\phi_0\rangle - \gamma^{-1}|j\rangle|1\rangle|\phi_1\rangle. \quad (3.10)$$

The operator R'_2 acts as identity on the states

$$|0\rangle|0\rangle|\phi_0\rangle \quad \text{and} \quad |j-1\rangle|1\rangle|\phi_1\rangle - \gamma|j\rangle|0\rangle|\phi_0\rangle \quad (3.11)$$

for all positive integers j . It reflects the following states for all positive integers j :

$$|j-1\rangle|1\rangle|\phi_1\rangle + \gamma^{-1}|j\rangle|0\rangle|\phi_0\rangle. \quad (3.12)$$

Proof. We prove the statements for R'_1 , as those for R'_2 are virtually identical. Since $|0\rangle|1\rangle|\phi_1\rangle$ has $|0\rangle|1\rangle$ in the first two registers, none of the controls gets activated, and so all gates act as the identity. For any $j > 0$, using $\gamma = \sqrt{p/(1-p)}$, and the fact that O_{ref} fixes $\sqrt{1-p}|0\rangle|\phi_0\rangle + \sqrt{p}|1\rangle|\phi_1\rangle$,


$$\begin{aligned} |j-1\rangle|0\rangle|\phi_0\rangle + \gamma|j\rangle|1\rangle|\phi_1\rangle &\xrightarrow{+1} \frac{|j\rangle}{\sqrt{1-p}} \left(\sqrt{1-p}|0\rangle|\phi_0\rangle + \sqrt{p}|1\rangle|\phi_1\rangle \right) \\ &\xrightarrow{O_{\text{ref}}} \frac{|j\rangle}{\sqrt{1-p}} \left(\sqrt{1-p}|0\rangle|\phi_0\rangle + \sqrt{p}|1\rangle|\phi_1\rangle \right) \\ &\xrightarrow{-1} |j-1\rangle|0\rangle|\phi_0\rangle + \gamma|j\rangle|1\rangle|\phi_1\rangle. \end{aligned}$$

Similarly, using the fact that O_{ref} reflects $\sqrt{p}|0\rangle|\phi_0\rangle - \sqrt{1-p}|1\rangle|\phi_1\rangle$,

$$\begin{aligned} |j-1\rangle|0\rangle|\phi_0\rangle - \gamma^{-1}|j\rangle|1\rangle|\phi_1\rangle &\xrightarrow{+1} \frac{|j\rangle}{\sqrt{p}} \left(\sqrt{p}|0\rangle|\phi_0\rangle - \sqrt{1-p}|1\rangle|\phi_1\rangle \right) \\ &\xrightarrow{O_{\text{ref}}} -\frac{|j\rangle}{\sqrt{p}} \left(\sqrt{p}|0\rangle|\phi_0\rangle - \sqrt{1-p}|1\rangle|\phi_1\rangle \right) \\ &\xrightarrow{-1} -(|j-1\rangle|0\rangle|\phi_0\rangle - \gamma^{-1}|j\rangle|1\rangle|\phi_1\rangle). \quad \square \end{aligned}$$

The purifier $S' = R'_2 R'_1$ can be interpreted as a quantum walk on a graph consisting of two disjoint rays, as described in Figure 3.3. Concretely, in Lemma 3.5, we show $|0\rangle|0\rangle|\phi_0\rangle \xrightarrow{S'(O_{\text{ref}})} (-1)^r |0\rangle|0\rangle|\phi_0\rangle$, where $r = 0$ if $p < 1/2$ and $r = 1$ if $p > 1/2$; and a similar statement for $|0\rangle|0\rangle|\phi_1\rangle$ in Lemma 3.6; and finally, in Theorem 3.7, we combine these to show that $|0\rangle|\phi\rangle \xrightarrow{S'(O_{\text{ref}})} (-1)^r |0\rangle|\phi\rangle$, proving the main theorem of this section.

Figure 3.3



As shown in Lemmata 3.5 and 3.6, the purifier in this section can be seen as a quantum walk on two copies of the ray graph from Figure 1.4. Just as in Figure 1.4, the weight of the edge between vertices j and $j + 1$, in either copy, is γ^{2j} (not pictured). The top ray represents Lemma 3.5, and the bottom one Lemma 3.6. Below each edge, the state that encodes it in the quantum walk is written. The pattern repeats with period 4. All these states are pairwise orthogonal. Above the edge, the corresponding basis state of the register \mathcal{N}_0 or \mathcal{N}_1 from the proof of the corresponding lemma is given.

Lemma 3.5. *In the above notation, the transducer S' , given an input oracle O_{ref} satisfying (1.6), has the following transduction action:*

$$S'(O_{\text{ref}}): |0\rangle|0\rangle|\phi_0\rangle \rightsquigarrow \begin{cases} |0\rangle|0\rangle|\phi_0\rangle, & \text{if } p < \frac{1}{2}; \\ -|0\rangle|0\rangle|\phi_0\rangle, & \text{if } p > \frac{1}{2}. \end{cases} \quad (3.13)$$

The query complexity and the transduction complexity of S' are $L(S', O_{\text{ref}}, |0\rangle|0\rangle|\phi_0\rangle) = 1/(2\delta)$ and $W(S', O_{\text{ref}}, |0\rangle|0\rangle|\phi_0\rangle) = \mathcal{O}(1/\delta)$, where $\delta = |\frac{1}{2} - p|$.

Proof. We will construct an invariant subspace \mathcal{N}_0 of the operator $S'(O_{\text{ref}})$ that contains the vector $|0\rangle|0\rangle|\phi_0\rangle$. We will define an isometry from \mathcal{N} as defined in Section 3.1 onto \mathcal{N}_0 such that $|0\rangle_{\mathcal{N}}$ gets mapped into $|0\rangle|0\rangle|\phi_0\rangle$, and the action of $S'(O_{\text{ref}})$ in \mathcal{N}_0 is identical (under this isometry) to the action of $S(O_p)$ in \mathcal{N} . The result then follows from Theorem 3.2.

To describe the isometry we will define an orthonormal system of vectors $|j\rangle_{\mathcal{N}_0}$ such that the isometry maps $|j\rangle_{\mathcal{N}}$ into $|j\rangle_{\mathcal{N}_0}$ for all non-negative integers j . These vectors are defined by

$$\begin{aligned} |4i\rangle_{\mathcal{N}_0} &= |4i\rangle|0\rangle|\phi_0\rangle, \\ |4i+1\rangle_{\mathcal{N}_0} &= |4i+1\rangle|1\rangle|\phi_1\rangle, \\ |4i+2\rangle_{\mathcal{N}_0} &= -|4i+2\rangle|0\rangle|\phi_0\rangle, \\ |4i+3\rangle_{\mathcal{N}_0} &= -|4i+3\rangle|1\rangle|\phi_1\rangle, \end{aligned}$$

as i ranges over non-negative integers (see Figure 3.3, the top ray). It is easy to check, using (3.9) and (3.10), that, for all positive odd values of j , R'_1 acts as identity on the states

$$|j-1\rangle_{\mathcal{N}_0} + \gamma|j\rangle_{\mathcal{N}_0} \quad (3.14)$$

and reflects the states

$$|j-1\rangle_{\mathcal{N}_0} - \gamma^{-1}|j\rangle_{\mathcal{N}_0}. \quad (3.15)$$

Similarly, using (3.11) and (3.12), we get that R'_2 acts as identity on $|0\rangle_{\mathcal{N}_0}$, and, for all positive even values of j , R'_2 acts as identity on the states in (3.14) and reflects the states in (3.15).

Comparing this with Claim 3.1, we get that the action of $R'_2R'_1$ in \mathcal{N}_0 is identical to that of R_2R_1 in \mathcal{N} , which proves the lemma. \square

Lemma 3.6. *In the above notation, the transducer S' , given an input oracle O_{ref} satisfying (1.6), has the following transduction action:*

$$S'(O_{\text{ref}}): |0\rangle|1\rangle|\phi_1\rangle \rightsquigarrow \begin{cases} |0\rangle|1\rangle|\phi_1\rangle, & \text{if } p < \frac{1}{2}; \\ -|0\rangle|1\rangle|\phi_1\rangle, & \text{if } p > \frac{1}{2}. \end{cases} \quad (3.16)$$

The query complexity and the transduction complexity of S' are $L(S', O_{\text{ref}}, |0\rangle|1\rangle|\phi_1\rangle) = 1/(2\delta)$ and $W(S', O_{\text{ref}}, |0\rangle|1\rangle|\phi_1\rangle) = \mathcal{O}(1/\delta)$, where $\delta = |\frac{1}{2} - p|$.

Proof. The idea is similar to Lemma 3.5. This time we define an invariant subspace \mathcal{N}_1 and an isometry from \mathcal{N} onto \mathcal{N}_1 that maps $|j\rangle_{\mathcal{N}}$ into $|j\rangle_{\mathcal{N}_1}$ for all non-negative integers j . The vectors are defined by

$$\begin{aligned} |4i\rangle_{\mathcal{N}_1} &= |4i\rangle|1\rangle|\phi_1\rangle, \\ |4i+1\rangle_{\mathcal{N}_1} &= -|4i+1\rangle|0\rangle|\phi_0\rangle, \\ |4i+2\rangle_{\mathcal{N}_1} &= -|4i+2\rangle|1\rangle|\phi_1\rangle, \\ |4i+3\rangle_{\mathcal{N}_1} &= |4i+3\rangle|0\rangle|\phi_0\rangle, \end{aligned}$$

as i ranges over non-negative integers (see Figure 3.3, the bottom ray). Using (3.11) and (3.12), it is not hard to check that, for all positive odd values of j , R'_2 acts as identity on the states

$$|j-1\rangle_{\mathcal{N}_1} + \gamma|j\rangle_{\mathcal{N}_1} \quad (3.17)$$

and reflects the states

$$|j-1\rangle_{\mathcal{N}_1} - \gamma^{-1}|j\rangle_{\mathcal{N}_1}. \quad (3.18)$$

Similarly, we get that R'_1 acts as identity on $|0\rangle_{\mathcal{N}_1}$ and, for all positive even j , R'_1 acts as identity on the states in (3.17) and reflects the states in (3.18).

Comparing this with Claim 3.1, we get that the action of $R'_2R'_1$ in \mathcal{N}_1 is identical to that of R_1R_2 in \mathcal{N} . The lemma follows by Remark 3.3. \square

Theorem 3.7. *The transducer S' defined in (3.8), given an input oracle O_{ref} satisfying (1.6), has the following transduction action:*

$$S'(O_{\text{ref}}): |\phi\rangle \rightsquigarrow \begin{cases} |\phi\rangle, & \text{if } p < \frac{1}{2}; \\ -|\phi\rangle, & \text{if } p > \frac{1}{2}; \end{cases} \quad (3.19)$$

for any normalised $|\phi\rangle \in \text{span}\{|0\rangle|\phi_0\rangle, |1\rangle|\phi_1\rangle\}$ (including the original vector $|\phi\rangle$ of (1.1)). The query complexity of the transducer is $L(S', O_{\text{ref}}, \phi) = 1/(2\delta)$, and the transduction complexity is $W(S', O_{\text{ref}}, \phi) = \mathcal{O}(1/\delta)$, where $\delta = |\frac{1}{2} - p|$.

The transducer uses infinite space. It can be put into canonical form (still using infinite space) so that its query complexity stays the same, and its transduction complexity remains $\mathcal{O}(1/\delta)$.

Proof. Recall that we identify $|0\rangle \otimes \mathcal{A} \otimes \mathcal{W}$, which is the public space of the transducer S' , with $\mathcal{A} \otimes \mathcal{W}$. Let $r = 0$ if $p < 1/2$; and $r = 1$ if $p > 1/2$. By [Lemma 3.5](#), we have

$$|0\rangle|\phi_0\rangle \xrightarrow{S'(O_{\text{ref}})} (-1)^r |0\rangle|\phi_0\rangle. \quad (3.20)$$

Let q_0 be the corresponding total query state $q_0 = q(S', O_{\text{ref}}, |0\rangle|\phi_0\rangle)$. By the proof, we have that S' acts in \mathcal{N}_0 , hence, $q_0 \in \mathbb{C}^2 \otimes \mathcal{N}_0$ as S' makes two queries. Also, $\|q_0\|^2 = 1/(2\delta)$. Similarly, by [Lemma 3.6](#),

$$|1\rangle|\phi_1\rangle \xrightarrow{S'(O_{\text{ref}})} (-1)^r |1\rangle|\phi_1\rangle. \quad (3.21)$$

Let $q_1 = q(S', O_{\text{ref}}, |1\rangle|\phi_1\rangle)$. Similarly as above, $q_1 \in \mathbb{C}^2 \otimes \mathcal{N}_1$ and $\|q_1\|^2 = 1/(2\delta)$.

By linearity [Claim 2.8](#), for all $\alpha, \beta \in \mathbb{C}$, we get from (3.20) and (3.21):

$$\alpha|0\rangle|\phi_0\rangle + \beta|1\rangle|\phi_1\rangle \xrightarrow{S'(O_{\text{ref}})} (-1)^r (\alpha|0\rangle|\phi_0\rangle + \beta|1\rangle|\phi_1\rangle), \quad (3.22)$$

and the corresponding total query state is $\alpha q_0 + \beta q_1$. Since q_0 and q_1 belong to orthogonal subspaces, we have that the query complexity of this transduction is

$$\|\alpha q_0 + \beta q_1\|^2 = |\alpha|^2 \|q_0\|^2 + |\beta|^2 \|q_1\|^2 = 1/(2\delta)$$

assuming $|\alpha|^2 + |\beta|^2 = 1$. The bound on the transduction complexity is proven similarly. The final statement about conversion into the canonical form follows directly from [Theorem 2.11](#). \square

3.3 Purifier for the State-Generating Input Oracle

In [Theorem 3.7](#), we established that the purifier S' exactly encodes the value of the Boolean function into the phase, given one copy of the state $|\phi\rangle$. It is good enough for applications. For completeness though, we obtain a purifier in the usual form similar to the one stated in Eqs. (1.1) and (1.2).

Corollary 3.8. *There exists a transducer S_{pur} with bidirectional access to an input oracle $O: \mathcal{A} \otimes \mathcal{W} \rightarrow \mathcal{A} \otimes \mathcal{W}$ as in (1.1), which performs the following transduction exactly:*

$$S_{\text{pur}}(O \oplus O^*): |0\rangle \rightsquigarrow \begin{cases} |0\rangle, & \text{if } p < \frac{1}{2}; \\ |1\rangle, & \text{if } p > \frac{1}{2}. \end{cases} \quad (3.23)$$

The query complexity $L(S_{\text{pur}}, O \oplus O^, |0\rangle) = 1 + 1/(2\delta)$, where $\delta = |\frac{1}{2} - p|$, and the transduction complexity is $W(S_{\text{pur}}, O \oplus O^*, |0\rangle) = \mathcal{O}(1/\delta)$. The transducer uses infinite space.*

Proof. In the proof, we again denote $r = 0$ if $p < 1/2$, and $r = 1$ if $p > 1/2$.

Let us first assume temporarily that we have algorithms (not transducers) performing the required actions, and we will explain how to combine them. The transducer result then follows by using compositions of transducers instead of algorithms.

Let $U' = S'(O_{\text{ref}})_{\mathcal{H}}$, which is given by (3.19). Assume we have an algorithm (not transducer) performing U' . Then, it is straightforward to obtain the transformation as in (3.23) as follows. (In the transformation below, the first register is the qubit used

in (3.23), and the second register is $\mathcal{A} \otimes \mathcal{W}$ as in (1.1).)

$$\begin{aligned}
|0\rangle|0\rangle &\longmapsto \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|0\rangle) && \text{Hadamard on the first register} \\
&\longmapsto \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|\phi\rangle) && \text{controlled application of } O \\
&\longmapsto \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + (-1)^r|1\rangle|\phi\rangle) && \text{controlled application of } U' \\
&\longmapsto \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + (-1)^r|1\rangle|0\rangle) && \text{controlled application of } O^* \\
&\longmapsto |r\rangle|0\rangle && \text{Hadamard on the first register}
\end{aligned} \tag{3.24}$$

One problem, though, is that S' uses the input oracle O_{ref} , not O . But it is simple to simulate the former with the latter, as by (1.5) we have $O_{\text{ref}} = O\text{Ref}_{|0\rangle|0\rangle}O^*$.

Let $A = A(O \oplus O^* \oplus O')$ be the algorithm in (3.24), where O' is the placeholder for U' . We have

$$L^{(0)}(A, O \oplus O^* \oplus U', |0\rangle|0\rangle) = \frac{1}{2} + \frac{1}{2} = 1, \tag{3.25}$$

where the superscript (0) corresponds to the input oracle $O \oplus O^*$. This is because $O \oplus O^*$ is called twice, each time on a vector of norm $1/\sqrt{2}$. Also, we have

$$q^{(1)}(A, O \oplus O^* \oplus U', |0\rangle|0\rangle) = \frac{(-1)^r \phi}{\sqrt{2}}, \tag{3.26}$$

where the superscript (1) corresponds to the input oracle O' . By Corollary 2.12, there exists a canonical transducer S_A whose transduction action is identical to the action of A , and that has the same total query state as A . Its transduction complexity is $\mathcal{O}(1)$.

Similarly, let $A_{\text{ref}} = A_{\text{ref}}(O \oplus O^*)$ be the algorithm from (1.5) whose action is equal to O_{ref} . It satisfies

$$L(A_{\text{ref}}, O \oplus O^*, \xi) = 2 \tag{3.27}$$

for every normalised $\xi \in \mathcal{E} \otimes \mathcal{A} \otimes \mathcal{W}$, where we use the convention from (2.22). Again, by Corollary 2.12 there exists a canonical transducer S_{ref} whose transduction action is equal to the action of A_{ref} and that has query and transduction complexity equal to 2 on every normalised initial state.

We get the required transducer S_{pur} as $S_A \circ S' \circ S_{\text{ref}}$. Using Theorem 2.17 to evaluate query complexity of composed transducers, we have

$$\begin{aligned}
L(S_{\text{pur}}, O \oplus O^*, |0\rangle|0\rangle) &= L^{(0)}(S_A, O \oplus O^* \oplus U', |0\rangle|0\rangle) \\
&\quad + L(S' \circ S_{\text{ref}}, O \oplus O^*, \frac{(-1)^r \phi}{\sqrt{2}}) && \text{by (3.26)} \\
&= 1 + \frac{1}{2}L(S' \circ S_{\text{ref}}, O \oplus O^*, \phi) && \text{by (3.25) and (2.15)} \\
&= 1 + \frac{1}{2}L(S_{\text{ref}}, O \oplus O^*, q(S', O_{\text{ref}}, \phi)) && \text{since } S' \text{ only queries } O_{\text{ref}} \\
&= 1 + \frac{1}{2} \cdot 2 \cdot \|q(S', O_{\text{ref}}, \phi)\|^2 && \text{by (3.27) and (2.15)} \\
&= 1 + \frac{1}{2} \cdot 2 \cdot \frac{1}{2\delta} = 1 + \frac{1}{2\delta} && \text{by Theorem 3.7.}
\end{aligned}$$

In a similar way, we get that the transduction complexity of S_{pur} is $\mathcal{O}(1/\delta)$. \square

4 Finite-Dimensional Implementations

The purifiers of [Theorems 3.2](#) and [3.7](#) cannot be implemented directly because they use registers (\mathcal{N} and \mathcal{J} , respectively) of infinite dimension. In this section, we consider two different variants of restricting the purifier to finite-dimensional space.

The first one is quite obvious: we replace the infinite-dimensional register with a D -dimensional register storing values between 0 and $D - 1$. This introduces small perturbation, but the overall structure of the algorithm is preserved. In particular, it can be implemented time-efficiently using the circuits from [Figures 3.1](#) and [3.2](#).

The second one is restricted to the special case of the state-conversion problem defined in [Section 2.6](#), where we have to implement the transformation $\xi_x \mapsto \tau_x$ for a finite number of pairs. In this case, we just restrict the infinite-dimensional space of the transducer to the finite-dimensional space spanned by the vectors used by these transformations.

4.1 Time-Efficient Implementation

As mentioned above, to obtain a time-efficient implementation, we replace the infinite-dimensional register with a D -dimensional register \mathbb{C}^D with the basis $|0\rangle, \dots, |D - 1\rangle$. We will call it the truncated version (at depth D) of the purifier. We assume $D \geq 4$ is a power of 2, which makes implementation easier.

The truncated version $S^{[D]} = S^{[D]}(O_{\text{ref}})$ of the purifier S from [Section 3.1](#) is similarly to [\(3.3\)](#) defined by

$$S^{[D]} = R_2^{[D]} R_1^{[D]},$$

where $R_1^{[D]}$ and $R_2^{[D]}$ are defined as R_1 and R_2 in [Figure 3.1](#), but where the register \mathcal{N} stores an integer between 0 and $D - 1$ and the operations are performed modulo D .

The truncated version $S'^{[D]} = S'^{[D]}(O_{\text{ref}})$ of the purifier S' from [Section 3.2](#) is defined similarly to [\(3.8\)](#) by

$$S'^{[D]} = R_2'^{[D]} R_1'^{[D]},$$

where $R_1'^{[D]}$ and $R_2'^{[D]}$ are defined as R_1' and R_2' in [Figure 3.2](#) with the same replacement for the register \mathcal{J} . The following claim is obvious.

Claim 4.1. *Both truncated purifiers $S^{[D]}$ and $S'^{[D]}$ use at most $\log D$ qubits besides the space used by the input oracle O_{ref} . Iteration times $T(S^{[D]})$ and $T(S'^{[D]})$ are $\mathcal{O}(\log D)$.*

From now on, we will concentrate on $S^{[D]}$, as the arguments for $S'^{[D]}$ are similar. For the former, we have the following simple claim.

Claim 4.2. *The actions of $R_1^{[D]}$ and R_1 are the same on all vectors in the span of $\{|0\rangle, \dots, |D - 1\rangle\}$. The actions of $R_2^{[D]}$ and R_2 are the same for all vectors in the span of $\{|0\rangle, \dots, |D - 2\rangle\}$. Also, $R_2^{[D]}|D - 1\rangle = |D - 1\rangle$ (contrary to R_2).*

Proof. Indeed, repeating the argument from the proof of [Claim 3.1](#), $R_1^{[D]}$ and R_1 act in the same way on all u_j and u_j^\perp for $j = 1, 3, 5, \dots, D - 1$, which are defined in [\(3.4\)](#). These vectors span the subspace $\text{span}\{|0\rangle, \dots, |D - 1\rangle\}$, which proves the claim for $R_1^{[D]}$.

Using similar reasoning, $R_2^{[D]}$ and R_2 act in the same way on $|0\rangle$ and all u_j and u_j^\perp for $j = 2, 4, 6, \dots, D - 2$. These vectors span the subspace $\text{span}\{|0\rangle, \dots, |D - 2\rangle\}$. The equality $R_2^{[D]}|D - 1\rangle = |D - 1\rangle$ is proven directly. \square

There are two ways to show that the truncated purifiers still work if D is large enough. The first one is based on the idea that the quantum walk will not get to the vertex D if the number of iterations is small enough.

Proposition 4.3. *Consider either the purifier S from [Section 3.1](#) or the purifier S' from [Section 3.2](#). Consider also the algorithm implementing the transduction action of the purifier, either from [Theorem 2.6](#) or [2.10](#). Then, if $D > 2K$, the action of the algorithm stays the same if we replace the actual purifier with its truncated version at depth D .*

Proof. We consider the transducer S , the case of S' being similar. We also consider the algorithm from [Theorem 2.6](#) because we gave its proof in [Section 2.3](#). The argument for the second algorithm is similar.

Recall that S acts in the space \mathcal{N} , its public space \mathcal{H} is spanned by $|0\rangle_{\mathcal{N}}$, and its private space \mathcal{L} by $|j\rangle_{\mathcal{N}}$ with integer $j \geq 1$. We consider the execution of the algorithm from [Theorem 2.6](#). Note that we *do not* apply any perturbations as in the proof of the said theorem.

Let ζ_t be the content of the register \mathcal{L} in this algorithm after the t -th application of S . We claim that

$$\zeta_t \in \text{span}\{|1\rangle_{\mathcal{N}}, \dots, |2t\rangle_{\mathcal{N}}\}. \quad (4.1)$$

We prove this by induction on t . For the base case, $\zeta_0 = 0$ because the algorithm starts in the state [\(2.9\)](#).

The algorithm repeatedly couples the content of the register \mathcal{L} with a fresh copy of ξ/\sqrt{K} , where ξ is the initial state. In the case of the purifier S , the initial state $\xi = |0\rangle$. For the inductive step, assume that [\(4.1\)](#) holds for a particular value of t . Consider the next application of S . Since R_1 and R_2 perform local reflections on the infinite ray in [Figure 1.4](#) (more formally, by [Claim 3.1](#)), we have that

$$R_1 \left(\frac{|0\rangle}{\sqrt{K}} \oplus \zeta_t \right) \in \text{span}\{|0\rangle_{\mathcal{N}}, \dots, |2t+1\rangle_{\mathcal{N}}\} \quad (4.2)$$

and

$$S \left(\frac{|0\rangle}{\sqrt{K}} \oplus \zeta_t \right) = R_2 R_1 \left(\frac{|0\rangle}{\sqrt{K}} \oplus \zeta_t \right) \in \text{span}\{|0\rangle_{\mathcal{N}}, \dots, |2t+2\rangle_{\mathcal{N}}\}. \quad (4.3)$$

This proves [\(4.1\)](#) with t increased by 1.

By [Claim 4.2](#), for all $t < D/2 - 1$, the action of R_1 in [\(4.2\)](#) is identical to the action of $R_1^{[D]}$, and the action of R_2 in [\(4.3\)](#) is identical to that of $R_2^{[D]}$. This proves the proposition. \square

A similar proof goes through also if the transducer S in [Proposition 4.3](#) is obtained via functional composition as in [Section 2.7](#) (or any other form of composition from [\[11\]](#)), where some of the constituents are purifiers S or S' .

The second way of justifying transition to the truncated purifier is based on perturbations as defined in [Section 2.8](#). [Theorem 1.7](#) from the introduction is a corollary of this result.

Theorem 4.4. *The statements of [Theorems 3.2](#) and [3.7](#) hold after replacing S and S' by their versions truncated at depth D if we allow for a perturbation of size $2(1-\delta)^{D-1}$ in the transducer, and claim that the query complexity is smaller than $1/(2\delta)$ (instead of being exactly equal).*

Proof. We start with the purifier S from [Section 3.1](#), and we mimic the proof of [Theorem 3.2](#) using [Claim 4.2](#). For $p < \frac{1}{2}$, we have the following version of [\(3.6\)](#), with no perturbations

$$\begin{aligned} \sum_{j=0}^{D-1} \gamma^j |j\rangle &= \sum_{i=0}^{D/2-1} \gamma^{2i} (|2i\rangle + \gamma |2i+1\rangle) \\ \xrightarrow{R_1^{[D]}} \sum_{j=0}^{D-1} \gamma^j |j\rangle &= |0\rangle + \sum_{i=0}^{D/2-2} \gamma^{2i+1} (|2i+1\rangle + \gamma |2i+2\rangle) + \gamma^{D-1} |D-1\rangle \\ \xrightarrow{R_2^{[D]}} \sum_{j=0}^{D-1} \gamma^j |j\rangle. \end{aligned}$$

For $p > \frac{1}{2}$, we have the following version of [\(3.7\)](#):

$$\begin{aligned} \sum_{j=0}^{D-1} (-\gamma)^{-j} |j\rangle &= \sum_{i=0}^{D/2-1} \gamma^{-2i} (|2i\rangle - \gamma^{-1} |2i+1\rangle) \\ \xrightarrow{R_1^{[D]}} - \sum_{j=0}^{D-1} (-\gamma)^{-j} |j\rangle &= -|0\rangle + \sum_{i=0}^{D/2-2} \gamma^{-2i-1} (|2i+1\rangle - \gamma^{-1} |2i+2\rangle) + \gamma^{-D+1} |D-1\rangle \\ \xrightarrow{R_2^{[D]}} -|0\rangle - \sum_{i=0}^{D/2-2} \gamma^{-2i-1} (|2i+1\rangle - \gamma^{-1} |2i+2\rangle) &+ \gamma^{-D+1} |D-1\rangle \\ \approx -|0\rangle - \sum_{i=0}^{D/2-2} \gamma^{-2i-1} (|2i+1\rangle - \gamma^{-1} |2i+2\rangle) &- \gamma^{-D+1} |D-1\rangle \\ = -|0\rangle + \sum_{j=1}^{D-1} (-\gamma)^{-j} |j\rangle, \end{aligned}$$

where the “ \approx ” represents a perturbation (in the sense of [Section 2.2](#)) of size

$$2\gamma^{-D+1} = 2 \left(\frac{1-p}{p} \right)^{\frac{D-1}{2}} = 2 \left(\frac{\frac{1}{2}-\delta}{\frac{1}{2}+\delta} \right)^{\frac{D-1}{2}} \leq 2(1-2\delta)^{\frac{D-1}{2}} \leq 2(1-\delta)^{D-1}.$$

The rest of the proof of [Theorem 3.2](#) stays the same except that the query and transduction complexities in both cases become slightly smaller.

The proof of [Theorem 3.7](#) follows the logic of [Section 3.2](#). We prove that [\(3.13\)](#) holds with S' replaced by $S'^{[D]}$ and perturbation $2(1-\delta)^{D-1}$ using a similar reduction as in [Lemma 3.5](#) but to the truncated and perturbed version of the purifier S . Similarly, we show the same for [\(3.16\)](#).

In the proof of [Theorem 3.7](#), we use the above two equations to derive [\(3.20\)](#) and [\(3.21\)](#) with the same modifications. In [\(3.22\)](#), we use that perturbations in [\(3.20\)](#) and [\(3.21\)](#) act in orthogonal spaces \mathcal{N}_0 and \mathcal{N}_1 , to get that the total perturbation is still of size $2(1-\delta)^{D-1}$. \square

4.2 Query-Efficient Implementation

As explained in [Section 2.6](#), in terms of query complexity, transducers have close connection to the adversary bound for state conversion. The state conversion problem is usually *finite*.

This means that it is given by pairs $\xi_x \mapsto \tau_x$ in \mathcal{H} and unitaries $O_x: \mathcal{M} \rightarrow \mathcal{M}$ where x runs through a *finite* set of labels X , and the spaces \mathcal{H} and \mathcal{M} are *finite-dimensional*. For instance, one can consider evaluation of a Boolean function, where x runs through the finite set $\{0, 1\}^n$ for some n . This motivates the following general result, which can be applied to all transducers in this paper.

Theorem 4.5. *Consider a state conversion problem given by pairs $\xi_x \mapsto \tau_x$ in \mathcal{H} and input oracles $O_x: \mathcal{M} \rightarrow \mathcal{M}$ as x ranges over a finite set X , and \mathcal{H} and \mathcal{M} are finite-dimensional. Let $S = S(O)$ be a transducer with infinite-dimensional private space that solves this problem, i.e., $\xi_x \xrightarrow{S(O_x)} \tau_x$ for every $x \in X$. Then, there exists a canonical transducer $S' = S'(O)$ in a finite-dimensional space such that $\xi_x \xrightarrow{S'(O_x)} \tau_x$ and $L(S', O_x, \xi_x) = L(S, O_x, \xi_x)$ for every $x \in X$.*

Proof. By [Theorem 2.15](#), the total query states v_x of S give a feasible solution to the adversary bound [\(2.18\)](#). One complication is that $v_x \in \mathcal{H}^\uparrow \otimes \mathcal{M}$ for some infinite-dimensional \mathcal{H}^\uparrow . Since the space \mathcal{M} has finite dimension, we have, by Schmidt decomposition, that $v_x \in \mathcal{H}_x \otimes \mathcal{M}$ for some finite-dimensional subspace $\mathcal{H}_x \subset \mathcal{H}^\uparrow$. Let $\mathcal{H}^\downarrow = \text{span}_{x \in X} \mathcal{H}_x$. It is a finite-dimensional subspace of \mathcal{H}^\uparrow and $v_x \in \mathcal{H}^\downarrow \otimes \mathcal{M}$ for all $x \in X$. By [Theorem 2.15](#) again, this implies existence of the transducer S' that uses finite-dimensional space. \square

Applying this to [Corollary 3.8](#), we get the following result.

Corollary 4.6. *Let \mathcal{A} be a qubit, and \mathcal{W} a finite-dimensional subspace. Consider a collection of states similar to the ones in [\(1.1\)](#)*

$$\phi_x = \sqrt{1 - p_x} |0\rangle_{\mathcal{A}} |\phi_{x,0}\rangle_{\mathcal{W}} + \sqrt{p_x} |1\rangle_{\mathcal{A}} |\phi_{x,1}\rangle_{\mathcal{W}}, \quad (4.4)$$

where x ranges over a finite set X , $\phi_{x,0}$ and $\phi_{x,1}$ are normalised, and $p_x \neq 1/2$. Let, for each $x \in X$, O_x be a unitary such that $|0\rangle \mapsto |\phi_x\rangle$ in $\mathcal{A} \otimes \mathcal{W}$. Then, there exists a finite-dimensional purifier S_{pur} such that

$$S_{\text{pur}}(O_x \oplus O_x^*): |0\rangle \rightsquigarrow \begin{cases} |0\rangle, & \text{if } p_x < \frac{1}{2}; \\ |1\rangle, & \text{if } p_x > \frac{1}{2} \end{cases}$$

with query complexity $1 + 1/(2\delta_x)$, where $\delta_x = \left| \frac{1}{2} - p_x \right|$.

Note that the construction in the proof of [Theorem 4.5](#) requires restriction to the subspace \mathcal{H}^\downarrow , which means that the corresponding transducer is *not* time-efficient. On the other hand, it yields *exact* purification. The new purifier has query complexity $\mathcal{O}(1/\delta)$, whereas the one from [\[6\]](#) has complexity $\mathcal{O}(1/\delta^2)$, see [Table 1.2](#).

This purifier can be used to convert bounded-error algorithms into exact dual adversary in the following way [\[6\]](#). Let A be a quantum query algorithm computing a function $f: X \rightarrow \{0, 1\}$ with bounded-error, and let O_x be the action of A on the input string $x \in X$. Then, composing it with the purifier S_{pur} from [Corollary 4.6](#), we get a *finite-dimensional* transducer that evaluates the function f *exactly* and whose query complexity is equal to the Las Vegas query complexity of A up to a constant factor (dependent on the gap δ). The latter transducer also yields the corresponding dual adversary via [Theorem 2.15](#).

5 Lower Bound

In this section, we show that query complexities of the purifiers S and S' from [Sections 3.1](#) and [3.2](#) are optimal. We prove this via an adversary lower bound for a very simple state conversion problem in the sense of [Definition 2.13](#). Assume $0 < \delta < 1/2$ is fixed. Let $X = \{0, 1\}$, $\mathcal{M} = \mathbb{C}^2$, and $\mathcal{H} = \mathbb{C}$. Consider two input oracles $O_0, O_1: \mathbb{C}^2 \rightarrow \mathbb{C}^2$ reflecting around the states

$$\varphi_0 = \sqrt{\frac{1}{2} + \delta} |0\rangle + \sqrt{\frac{1}{2} - \delta} |1\rangle \quad \text{and} \quad \varphi_1 = \sqrt{\frac{1}{2} - \delta} |0\rangle + \sqrt{\frac{1}{2} + \delta} |1\rangle, \quad (5.1)$$

respectively. I.e., $O_i = 2\varphi_i\varphi_i^* - I$. In notation of [Section 3.1](#), these are precisely the states $\varphi_{\frac{1}{2} \pm \delta}$ from [\(3.1\)](#), and the oracles are $O_{\frac{1}{2} \pm \delta}$, but we use 0 and 1 here for conciseness.

Let $\xi_0 = \xi_1 = \tau_0 = |0\rangle$, and $\tau_1 = -|0\rangle$. This state conversion problem is a restriction of the settings of [Theorem 3.2](#) for $p = \frac{1}{2} \pm \delta$, so the purifier of that theorem solves this problem in query complexity $1/(2\delta)$. This state conversion problem also falls into the settings of [Theorem 3.7](#) since the space \mathcal{M} is 2-dimensional, and the vector $|0\rangle$ belongs to $\text{span}\{|0\rangle|\phi_0\rangle, |1\rangle|\phi_1\rangle\}$ from the statement of the theorem. The transducer of the latter theorem also has query complexity $1/(2\delta)$. The next result shows that this cannot be improved.

Theorem 5.1. *For any $0 < \delta < 1/2$, any transducer $S = S(O)$ solving the above state conversion problem satisfies $\max\{L(S, O_0, \xi_0), L(S, O_1, \xi_1)\} \geq 1/(2\delta)$.*

Proof. We prove that any feasible solution v_0, v_1 to the corresponding dual adversary bound satisfies $\max\{\|v_0\|^2, \|v_1\|^2\} \geq 1/(2\delta)$. This proves the theorem by [Corollary 2.16\(a\)](#).

[Eq. \(2.17\)](#) with $x = 0$ and $y = 1$ read as

$$\langle \xi_0, \xi_1 \rangle - \langle \tau_0, \tau_1 \rangle = \langle v_0, I \otimes (I_{\mathcal{M}} - O_0^* O_1) v_1 \rangle.$$

This implies

$$\left| \langle \xi_0, \xi_1 \rangle - \langle \tau_0, \tau_1 \rangle \right| \leq \|v_0\| \cdot \|I \otimes (I_{\mathcal{M}} - O_0^* O_1)\| \cdot \|v_1\|. \quad (5.2)$$

The left-hand side of the above equation is 2, and for the right-hand side, we have

$$\|I \otimes (I_{\mathcal{M}} - O_0^* O_1)\| = \|I_{\mathcal{M}} - O_0^* O_1\| = \|O_0 - O_1\| = 2\|\varphi_0\varphi_0^* - \varphi_1\varphi_1^*\|.$$

Plugging this into [\(5.2\)](#), we get

$$\max\{\|v_0\|^2, \|v_1\|^2\} \geq \|v_0\| \cdot \|v_1\| \geq \frac{1}{\|\varphi_0\varphi_0^* - \varphi_1\varphi_1^*\|}. \quad (5.3)$$

It is easy to check that

$$\varphi_0\varphi_0^* - \varphi_1\varphi_1^* = \begin{pmatrix} 2\delta & 0 \\ 0 & -2\delta \end{pmatrix},$$

which gives the required lower bound by [\(5.3\)](#). \square

6 Error Reduction using Quantum Signal Processing

In this section, we describe an error reduction algorithm using quantum signal processing. To the best of our knowledge, this particular version has not been previously stated explicitly, however, very similar constructions have been considered in previous work – see [\[20, Section 3\]](#). We were particularly inspired by the discussion in [\[35\]](#). Note that we are

using usual Quantum Signal Processing [30] directly here like in [29], and not the more elaborate framework of Quantum Singular Value Transformation [20]. This streamlined construction uses no additional ancillary qubits, whereas Ref. [20] requires a small number, and our construction works with reflecting oracles of the form in (1.6), whereas Ref. [20] requires state-generating oracles of the form (1.1).

6.1 Quantum Signal Processing Preliminaries

Let x and y be real numbers satisfying

$$-1 \leq x, y \leq 1 \quad \text{and} \quad x^2 + y^2 = 1.$$

Define the following unitary transformation⁶

$$W = W(x, y) = \begin{pmatrix} x & y \\ y & -x \end{pmatrix}. \quad (6.1)$$

Let $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_k)$ be a sequence of unimodular complex numbers: $|\alpha_j| = 1$ for all j . Then, the *quantum signal processing* algorithm (QSP) corresponding to α is the following k -query quantum algorithm $U_\alpha = U_\alpha(W)$ with an input oracle $W: \mathbb{C}^2 \rightarrow \mathbb{C}^2$:

$$U_\alpha(W) = \begin{pmatrix} \alpha_k & \\ & -\alpha_k^* \end{pmatrix} W \begin{pmatrix} \alpha_{k-1} & \\ & -\alpha_{k-1}^* \end{pmatrix} W \cdots W \begin{pmatrix} \alpha_1 & \\ & -\alpha_1^* \end{pmatrix} W \begin{pmatrix} \alpha_0 & \\ & -\alpha_0^* \end{pmatrix}. \quad (6.2)$$

Here α^* stands for the complex conjugate of a complex number α .

Theorem 6.1 ([20]). *For any sequence $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_k)$ of unimodular complex numbers, there exist complex polynomials $P(x)$ and $Q(x)$ satisfying the following conditions:*

$$\deg P \leq k, \quad \deg Q \leq k-1, \quad (6.3)$$

$$P(-x) = (-1)^k P(x), \quad Q(-x) = (-1)^{k-1} Q(x), \quad (6.4)$$

$$P(x)P^*(x) + Q(x)Q^*(x)(1-x^2) = 1 \quad (6.5)$$

and such that

$$U_\alpha(W(x, y)) = \begin{pmatrix} P(x) & yQ^*(x) \\ yQ(x) & -P^*(x) \end{pmatrix}. \quad (6.6)$$

And vice versa: for every pair of complex polynomials $P(x)$ and $Q(x)$ and non-negative integer k satisfying (6.3), (6.4) and (6.5), there exists a sequence $\alpha = (\alpha_0, \alpha_1, \dots, \alpha_k)$ of unimodular complex numbers such that (6.6) holds.

Theorem 6.2 ([20]). *Let R be a real polynomial and k a positive integer satisfying*

$$\deg R \leq k, \quad R(-x) = (-1)^k R(x),$$

and

$$|R(x)| \leq 1 \quad \text{for all} \quad -1 \leq x \leq 1.$$

Then, there exist complex polynomials P and Q satisfying (6.3)–(6.5) and such that $R(x) = \Re(P(x))$, where \Re stands for the real part.

Theorem 6.3 ([29]). *Let $0 < \varepsilon', \delta' < 1$ be parameters. There exists an odd real polynomial $R(x)$ of degree $\mathcal{O}(\frac{1}{\varepsilon'} \log \frac{1}{\varepsilon'})$ such that*

$$\begin{aligned} |R(x)| &\leq 1 && \text{for all} && -1 \leq x \leq 1; \\ R(x) &\geq 1 - \varepsilon' && \text{for all} && \delta' \leq x \leq 1; \quad \text{and} \\ R(x) &\leq -1 + \varepsilon' && \text{for all} && -1 \leq x \leq -\delta'. \end{aligned} \quad (6.7)$$

$$R(x) \leq -1 + \varepsilon' \quad \text{for all} \quad -1 \leq x \leq -\delta'. \quad (6.8)$$

⁶Traditionally, a more complex matrix $\begin{pmatrix} x & iy \\ iy & x \end{pmatrix}$ is used. The two are equivalent up to a simple change of basis.

6.2 Proof of Theorem 1.4

We will work in the two-dimensional subspace \mathcal{H} given by the basis vectors $|0\rangle_{\mathcal{H}} = |0\rangle_{\mathcal{A}}|\phi_0\rangle_{\mathcal{W}}$ and $|1\rangle_{\mathcal{H}} = |1\rangle_{\mathcal{A}}|\phi_1\rangle_{\mathcal{W}}$. By (1.6), the operator O_{ref} performs reflection about $\sqrt{1-p}|0\rangle_{\mathcal{H}} + \sqrt{p}|1\rangle_{\mathcal{H}}$ in this space, and is given by the matrix

$$O_{\text{ref}} = \begin{pmatrix} 1 - 2p & 2\sqrt{p(1-p)} \\ 2\sqrt{p(1-p)} & 2p - 1 \end{pmatrix}.$$

This matrix is in the form (6.1) with $x = 1 - 2p$ and $y = 2\sqrt{p(1-p)}$. Thus, $x \geq 2\delta$ if $p \leq \frac{1}{2} - \delta$, and $x \leq -2\delta$ if $p \geq \frac{1}{2} + \delta$.

Take the real polynomial R of Theorem 6.3 with $\delta' = 2\delta$ and ε' to be fixed later. Find the complex polynomials P and Q by Theorem 6.2 applied to R and $k = \deg R$. Plug them into Theorem 6.1 to get a sequence $\alpha = (\alpha_0, \dots, \alpha_k)$ such that (6.6) holds. We apply the Quantum Signal Processing algorithm U_α from (6.2) in \mathcal{H} by replacing W with O_{ref} and applying the intermediate operations $\begin{pmatrix} \alpha_j & \\ & -\alpha_j^* \end{pmatrix}$ to the register \mathcal{A} . After that, we apply the Z operator to \mathcal{A} . By Theorem 6.1, the resulting action of this algorithm on \mathcal{H} is

$$U(p) = \begin{pmatrix} P(x) & yQ^*(x) \\ -yQ(x) & P^*(x) \end{pmatrix}.$$

Let $P(x) = R(x) + iS(x)$ for a real polynomial S . Since the matrix $U(p)$ is unitary (alternatively, by (6.5)), we have

$$R(x)^2 + S(x)^2 + |yQ(x)|^2 = 1. \quad (6.9)$$

If $p \leq \frac{1}{2} - \delta$, then by (6.7), $1 - \varepsilon' \leq R(x) \leq 1$, and, by (6.9), $S(x)^2 + |yQ(x)|^2 \leq 1 - (1 - \varepsilon')^2 \leq 2\varepsilon'$. Using that the spectral norm is bounded by the Frobenius norm, we get

$$\begin{aligned} \|U(p) - I\|^2 &= \left\| \begin{pmatrix} R(x) - 1 + iS(x) & yQ^*(x) \\ -yQ(x) & R(x) - 1 - iS(x) \end{pmatrix} \right\|^2 \\ &\leq 2 \left((R(x) - 1)^2 + S(x)^2 + |yQ(x)|^2 \right) \\ &\leq 6\varepsilon'. \end{aligned}$$

Similarly, using (6.8), we have that if $p \geq \frac{1}{2} + \delta$, then

$$\|U(p) + I\| \leq \sqrt{6\varepsilon'}.$$

Letting $r = 0$ if $p \leq \frac{1}{2} - \delta$ and $r = 1$ if $p \geq \frac{1}{2} + \delta$, we obtain

$$\|U(p)|\phi\rangle - (-1)^r|\phi\rangle\| \leq \sqrt{6\varepsilon'}.$$

Therefore, we can take $\varepsilon' = \varepsilon^2/6$ to obtain the required precision of the algorithm.

The number of queries and time complexity of the algorithm is $\mathcal{O}(\deg R)$, which is $\mathcal{O}\left(\frac{1}{\delta'} \log \frac{1}{\varepsilon'}\right) = \mathcal{O}\left(\frac{1}{\delta} \log \frac{1}{\varepsilon}\right)$ as required.

7 Non-Boolean Case

In this section, we briefly describe the non-Boolean error reduction case, which we reduce to the Boolean case using the ‘‘Bernstein-Vazirani’’ trick from [27, Section 4]. In the non-Boolean case, we can always assume the output register $\mathcal{A} = (\mathbb{C}^2)^{\otimes m}$ is a tensor product

of m qubits. We denote basis elements of \mathcal{A} by bit-strings in the vector space \mathbb{F}_2^m over the field $\mathbb{F}_2 = \{0, 1\}$. For $a = (a_1, \dots, a_m)$ and $b = (b_1, \dots, b_m)$ in \mathbb{F}_2^m , we denote by

$$a \odot b = a_1 b_1 + \dots + a_m b_m \in \{0, 1\}$$

their inner product, where the operations are in \mathbb{F}_2 . The work register \mathcal{W} is still arbitrary. We assume the input oracle O_{ref} reflects about some state

$$\phi = \sum_{a \in \mathbb{F}_2^m} \sqrt{p_a} |a\rangle_{\mathcal{A}} |\phi_a\rangle_{\mathcal{W}} \quad (7.1)$$

for normalised ϕ_a . Moreover, we assume there exists a (unique) r such that $p_r \geq \frac{1}{2} + \delta$. The task is to perform the transformation

$$|0\rangle_{\mathcal{B}} |\phi\rangle_{\mathcal{AW}} \mapsto |r\rangle_{\mathcal{B}} |\phi\rangle_{\mathcal{AW}}, \quad (7.2)$$

where \mathcal{B} is a register isomorphic to \mathcal{A} .

Theorem 7.1. *Let R be a Boolean error reduction subroutine that, given the oracle as in (1.5), performs the transformation in (1.7) (possibly with imprecision ε). Then, there exists an algorithm that, given an input oracle O_{ref} reflecting about the state (7.1), implements the transformation in (7.2) (with the same imprecision ε). The algorithm makes one execution of R , uses one additional qubit and $\mathcal{O}(qm)$ additional elementary operations, where q is the query complexity of R .*

Proof. We use the following transformation, where C is a qubit:

$$T: |b\rangle_{\mathcal{B}} |c\rangle_{\mathcal{C}} |a\rangle_{\mathcal{A}} \mapsto |b\rangle_{\mathcal{B}} |c \oplus (a \odot b)\rangle_{\mathcal{C}} |a\rangle_{\mathcal{A}}.$$

It evaluates the inner product of a and b into the register \mathcal{C} , and can be efficiently implemented. We can decompose T as $T = \bigoplus_{b \in \mathbb{F}_2^m} T_b$, where $T_b: |c\rangle_{\mathcal{C}} |a\rangle_{\mathcal{A}} \mapsto |c \oplus (a \odot b)\rangle_{\mathcal{C}} |a\rangle_{\mathcal{A}}$.

Consider the following state

$$\begin{aligned} \phi'_b &= T_b |0\rangle_{\mathcal{C}} |\phi\rangle_{\mathcal{AW}} = \sum_{a \in \mathbb{F}_2^m} \sqrt{p_a} |a \odot b\rangle_{\mathcal{C}} |a\rangle_{\mathcal{A}} |\phi_a\rangle_{\mathcal{W}} \\ &= \sqrt{1 - p'_b} |0\rangle_{\mathcal{C}} |\phi'_{b,0}\rangle_{\mathcal{AW}} + \sqrt{p'_b} |1\rangle_{\mathcal{C}} |\phi'_{b,1}\rangle_{\mathcal{AW}}, \end{aligned}$$

where all $\phi'_{b,c}$ are normalised. The main idea is that

$$p'_b \geq \frac{1}{2} + \delta, \text{ if } r \odot b = 1; \quad \text{and} \quad p'_b \leq \frac{1}{2} - \delta, \text{ otherwise.} \quad (7.3)$$

Thus, the output of the error reduction subroutine R encodes $r \odot b$ into the phase. From that, we can obtain b using the Bernstein-Vazirani algorithm [12].

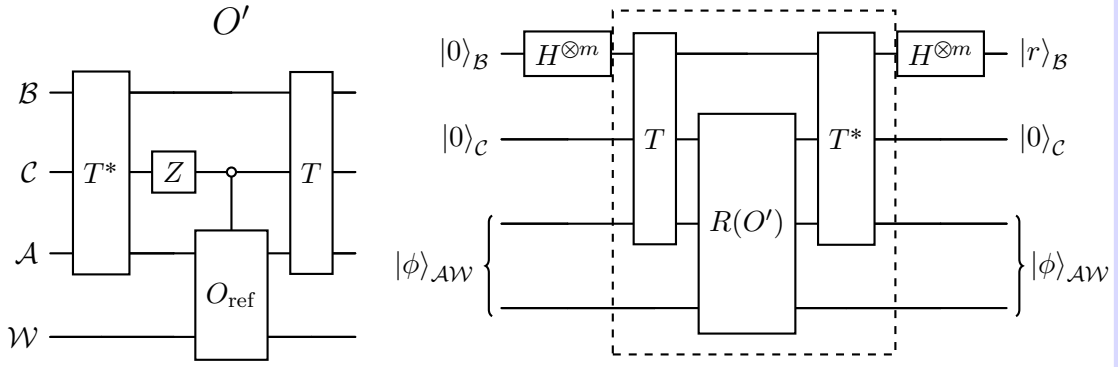
Let us fill in the technical details. We will first assume that R is exact. It is easy to see that the circuit O' in Figure 7.1 on the left implements the transformation

$$O' = \bigoplus_{b \in \mathbb{F}_2^m} O'_b, \quad (7.4)$$

where O'_b reflects about the state ϕ'_b in $\mathcal{C} \otimes \mathcal{A} \otimes \mathcal{W}$. Thus, using parallel composition (2.20) over the register \mathcal{B} , where all composed subroutines are equal to R , and taking (7.3) into consideration, we have

$$(I_{\mathcal{B}} \otimes R)(O'): \quad \frac{1}{\sqrt{2^m}} \bigoplus_{b \in \mathbb{F}_2^m} \phi'_b \mapsto \frac{1}{\sqrt{2^m}} \bigoplus_{b \in \mathbb{F}_2^m} (-1)^{r \odot b} \phi'_b, \quad (7.5)$$

Figure 7.1



Implementation of the oracle O' and the algorithm in [Theorem 7.1](#). Assuming R is exact, the operations in the dashed box effectively implement the Bernstein-Vazirani oracle for r the hidden string: $|b\rangle_{\mathcal{B}}|0\rangle_{\mathcal{C}}|\phi\rangle_{\mathcal{AW}} \mapsto (-1)^{r \odot b}|b\rangle_{\mathcal{B}}|0\rangle_{\mathcal{C}}|\phi\rangle_{\mathcal{AW}}$.

which implies

$$T^* \left[(I_{\mathcal{B}} \otimes R)(O') \right] T: \frac{1}{\sqrt{2^m}} \sum_{b \in \mathbb{F}_2^m} |b\rangle_{\mathcal{B}} |0\rangle_{\mathcal{C}} |\phi\rangle_{\mathcal{AW}} \mapsto \frac{1}{\sqrt{2^m}} \sum_{b \in \mathbb{F}_2^m} (-1)^{r \odot b} |b\rangle_{\mathcal{B}} |0\rangle_{\mathcal{C}} |\phi\rangle_{\mathcal{AW}}. \quad (7.6)$$

Using the same analysis as in the Bernstein-Vazirani algorithm [12], we get that the algorithm in [Figure 7.1](#) on the right implements the transformation in (7.2). The subcircuit in the dashed box implements the transformation in (7.6).

The above analysis assumed that R worked exactly. If R has imprecision ε , then the whole algorithm also has imprecision ε because the perturbations in different copies of R in $I_{\mathcal{B}} \otimes R$ occur in pairwise orthogonal subspaces. \square

This result can be combined with any Boolean error reduction or purification procedure. Combining with [Theorem 1.4](#), we obtain the following result.

Corollary 7.2. *Let $\varepsilon, \delta > 0$ be parameters. There exists a quantum algorithm that, given query access to the oracle O_{ref} reflecting about the state (7.1), ε -approximately performs the transformation in (7.2). The algorithm makes $\ell = \mathcal{O}(\frac{1}{\delta} \log \frac{1}{\varepsilon})$ queries to the oracle O_{ref} , uses one additional qubit, and has time complexity $\mathcal{O}(\ell + m)$.*

Similarly, we can combine this result with a purifier S' from [Theorem 3.7](#) to obtain the following result improving on [Theorem 13.3](#) of [8].

Corollary 7.3. *Let the registers \mathcal{A} and \mathcal{W} be fixed as at the beginning of this section. There exists a transducer $S'' = S''(O_{\text{ref}})$ with the following properties. If the input oracle O_{ref} reflects about the state $|\phi\rangle$ as in (7.1), S'' performs the transduction*

$$S''(O_{\text{ref}}): |0\rangle_{\mathcal{B}} |\phi\rangle_{\mathcal{AW}} \rightsquigarrow |r\rangle_{\mathcal{B}} |\phi\rangle_{\mathcal{AW}}.$$

The query complexity of the transduction is at most $1/(2\delta)$, where $\delta = p_r - 1/2$, its transduction complexity is $\mathcal{O}(1/\delta)$, and its iteration time complexity is $\mathcal{O}(m)$.

Proof. Let $O' = O'(O_{\text{ref}})$ and $A = A(R)$ be the algorithms in [Figure 7.1](#) on the left and on the right, respectively. We can write the algorithm of [Theorem 7.1](#) as the functional composition

$$A \circ R \circ O'. \quad (7.7)$$

We adhere to the same proof strategy as in [Corollary 3.8](#): We obtain the required transduction action by replacing R by S' , and A and O' by the corresponding transducers.

Let $S_O = S_O(O_{\text{ref}})$ and $S_A = S_A(R)$ be canonical transducers as in [Corollary 2.12](#) whose transduction action and query complexity is identical to the action of O' and A . Let $S' = S'(O')$ be the purifier from [Theorem 3.7](#) acting in $\mathcal{C} \otimes \mathcal{A} \otimes \mathcal{W}$. Then, the purifier $S'' = S''(O_{\text{ref}})$ can be written as the same functional composition as in [\(7.7\)](#):

$$S'' = S_A \circ S' \circ S_O.$$

The correctness of the transducer then follows immediately from the proof of [Theorem 7.1](#).

We analyse query complexity of the transduction using [Theorem 2.17](#), the analysis of transduction complexity being similar. Since S_A and S_O have the same query complexity as A and O' , respectively, it suffices to evaluate query complexities of the latter. The algorithm O' uses just one query, hence, $L(S_O, O_{\text{ref}}, \xi) \leq 1$ for every normalised initial state ξ . As $S_A \circ S'$ only accesses O_{ref} through O' , we have

$$\begin{aligned} L(S_A \circ S' \circ S_O, O_{\text{ref}}, |0\rangle_{\mathcal{B}}|0\rangle_{\mathcal{C}}|\phi\rangle_{\mathcal{AW}}) &= L\left(S_O, O_{\text{ref}}, q(S_A \circ S', O', |0\rangle_{\mathcal{B}}|0\rangle_{\mathcal{C}}|\phi\rangle_{\mathcal{AW}})\right) \\ &\leq L(S_A \circ S', O', |0\rangle_{\mathcal{B}}|0\rangle_{\mathcal{C}}|\phi\rangle_{\mathcal{AW}}), \end{aligned}$$

where we used [\(2.15\)](#) on the second step. Since A only accesses O' through its only query to R given in [\(7.5\)](#), we have

$$\begin{aligned} L(S_A \circ S', O', |0\rangle_{\mathcal{B}}|0\rangle_{\mathcal{C}}|\phi\rangle_{\mathcal{AW}}) &= L\left(S', O', \frac{1}{\sqrt{2^m}} \bigoplus_{b \in \mathbb{F}_2^m} \phi'_b\right) \\ &= \frac{1}{2^m} L\left(S', \bigoplus_{b \in \mathbb{F}_2^m} O'_b, \bigoplus_{b \in \mathbb{F}_2^m} \phi'_b\right) \quad \text{using (7.4) and (2.15)} \\ &= \frac{1}{2^m} \sum_{b \in \mathbb{F}_2^m} L(S', O'_b, \phi'_b) \quad \text{by (2.25)} \\ &\leq \frac{1}{2\delta} \quad \text{by Theorem 3.7 and (7.3).} \quad \square \end{aligned}$$

Time- and query-efficient implementations from [Sections 4.1](#) and [4.2](#) are applicable here as well. Note that the estimate $1/(2\delta)$ in [Corollary 7.3](#) is not tight. For instance, it is not necessary to execute R for $b = 0$, since $r \odot b = 0$ for every r . In addition, the inequalities in [\(7.3\)](#) will be strict for many values of b . We leave investigation of the precise value of the query complexity in the non-Boolean case as future work. Another interesting problem is error reduction when it is promised that there exists a unique r with somewhat large p_r , but whose value is less than $1/2$.

Acknowledgements AB is supported by the Latvian Quantum Initiative under European Union Recovery and Resilience Facility project no. 2.3.1.1.i.0/1/22/I/CFLA/001.

SJ is supported by NWO Klein project number OCENW.Klein.061; the project Divide & Quantum (with project number 1389.20.241) of the research programme NWA-ORC, which is (partly) financed by the Dutch Research Council (NWO); and by the European

Union (ERC, ASC-Q, 101040624). SJ is a CIFAR Fellow in the Quantum Information Science Program.

We thank the anonymous referees for comments that have substantially improved the presentation of this manuscript.

References

- [1] S. Aaronson, R. Kothari, W. Kretschmer, and J. Thaler. Quantum lower bounds for approximate counting via Laurent polynomials. In *Proc. of 35th IEEE CCC*, volume 169 of *LIPICs*, pages 7:1–7:47, 2020. [arXiv:1904.08914](#).
- [2] A. Ambainis. Quantum walk algorithm for element distinctness. *SIAM Journal on Computing*, 37(1):210–239, 2007. Earlier: *FOCS'04*, [arXiv:quant-ph/0311001](#).
- [3] A. Ambainis, A. M. Childs, B. W. Reichardt, R. Špalek, and S. Zhang. Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM Journal on Computing*, 39(6):2513–2530, 2010. Earlier: *FOCS'07*.
- [4] A. Ambainis, L. Magnin, M. Rötteler, and J. Roland. Symmetry-assisted adversaries for quantum state generation. In *Proc. of 26th IEEE CCC*, pages 167–177, 2011. [arXiv:1012.2112](#).
- [5] A. Belovs. Quantum walks and electric networks. [arXiv:1302.3143](#), 2013.
- [6] A. Belovs. Variations on quantum adversary. [arXiv:1504.06943](#), 2015.
- [7] A. Belovs. Global phase helps in quantum search: Yet another look at the welded tree problem. [arXiv:2404.19476](#), 2024.
- [8] A. Belovs, S. Jeffery, and D. Yolcu. Taming quantum time complexity. *Quantum*, 8:1444, 2024. [arXiv:2311.15873](#).
- [9] A. Belovs and T. Lee. Quantum algorithm for k -distinctness with prior knowledge on the input. [arXiv:1108.3022](#), 2011.
- [10] A. Belovs and A. Rosmanis. Tight quantum lower bound for approximate counting with quantum states. [arXiv:2002.06879](#), 2020.
- [11] A. Belovs and D. Yolcu. One-way ticket to Las Vegas and the quantum adversary. [arXiv:2301.02003](#), 2023.
- [12] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal on Computing*, 26(5):1411–1473, 1997. Earlier: *STOC'93*.
- [13] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information: A Millennium Volume*, volume 305 of *AMS Contemporary Mathematics Series*, pages 53–74, 2002. [arXiv:quant-ph/0005055](#).
- [14] H. Buhrman, R. Cleve, R. de Wolf, and C. Zalka. Bounds for small-error and zero-error quantum algorithms. In *Proc. of 40th IEEE FOCS*, pages 358–368, 1999. [arXiv:cs/9904019](#).
- [15] H. Buhrman and R. de Wolf. Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288:21–43, 2002.
- [16] A. M. Childs, R. Cleve, S. P. Jordan, and D. Yonge-Mallo. Discrete-query quantum algorithm for nand trees. *Theory OF Computing*, 5:119–123, 2009. [arXiv:quant-ph/0702160](#).
- [17] A. Cornelissen, S. Jeffery, M. Ozols, and A. Piedrafita. Span programs and quantum time complexity. In *Proc. of 45th MFCS*, pages 26:1–26:14, 2020. [arXiv:2005.01323](#).
- [18] P. G. Doyle and J. L. Snell. *Random walks and electric networks*, volume 22 of *Carus Mathematical Monographs*. MAA, 1984. [arXiv:math.PR/0001057](#).
- [19] E. Farhi, J. Goldstone, and S. Gutmann. A quantum algorithm for the Hamiltonian NAND tree. *Theory of Computing*, 4:169–190, 2008. [arXiv:quant-ph/0702144](#).
- [20] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proc. of 51st ACM STOC*, pages 193–204, 2019. [arXiv:1806.01838](#).

- [21] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proc. of 28th ACM STOC*, pages 212–219, 1996. [arXiv:quant-ph/9605043](#).
- [22] P. Høyer, T. Lee, and R. Špalek. Negative weights make adversaries stronger. In *Proc. of 39th ACM STOC*, pages 526–535, 2007. [arXiv:quant-ph/0611054](#).
- [23] T. Ito and S. Jeffery. Approximate span programs. *Algorithmica*, 81(6):2158–2195, 2019. [arXiv:1507.00432](#).
- [24] S. Jeffery. Span programs and quantum space complexity. *Theory of Computing*, 18(11):1–49, 2022. [arXiv:1908.04232](#).
- [25] S. Jeffery. SIGACT News Complexity Theory Column 123: Composing quantum algorithms. *ACM SIGACT News*, 55(4):49–69, December 2024.
- [26] S. Jeffery and F. Witteveen. QMA = QMA₁ with an infinite counter. [arXiv:2506.15551](#), 2025.
- [27] S. Jeffery and S. Zur. Multidimensional quantum walks, with application to k-distinctness. In *Proc. of 55th ACM STOC*, pages 1125–1130, 2023. [arXiv:2208.13492](#).
- [28] T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of state conversion. In *Proc. of 52nd IEEE FOCS*, pages 344–353, 2011. [arXiv:1011.3020](#).
- [29] G. H. Low and I. L. Chuang. Hamiltonian simulation by uniform spectral amplification. [arXiv:1707.05391](#), 2017.
- [30] G. H. Low, T. J. Yoder, and I. L. Chuang. Methodology of resonant equiangular composite quantum gates. *Physical Review X*, 6(4):041067, 2016. [arXiv:1603.03996](#).
- [31] C. Marriott and J. Watrous. Quantum Arthur–Merlin games. *Computational Complexity*, 14(2):122–152, 2005. Earlier: *CCC’04*, [arXiv:cs/0506068](#).
- [32] D. Nagaj, P. Wocjan, and Y. Zhang. Fast amplification of QMA. *Quantum Information and Computation*, 9(11-12):1053–1068, 2009. [arXiv:0904.1549](#).
- [33] A. Nayak and F. Wu. The quantum query complexity of approximating the median and related statistics. In *Proc. of 31st ACM STOC*, pages 384–393, 1999. [arXiv:quant-ph/9804066](#).
- [34] G. Pólya. Über eine aufgabe der wahrscheinlichkeitsrechnung betreffend die irrfahrt im straßennetz. *Mathematische Annalen*, 84(1):149–160, 1921.
- [35] P. Rall. Faster coherent quantum algorithms for phase, energy, and amplitude estimation. *Quantum*, 5:566, 2021. [arXiv:2103.09717](#).
- [36] B. W. Reichardt. Span programs and quantum query complexity: The general adversary bound is nearly tight for every Boolean function. In *Proc. of 50th IEEE FOCS*, pages 544–551, 2009. [arXiv:0904.2759](#).
- [37] B. W. Reichardt and R. Špalek. Span-program-based quantum algorithm for evaluating formulas. *Theory of Computing*, 8:291–319, 2012. Earlier: *STOC’08*, [arXiv:0710.2630](#).