

Bond University
Research Repository



A review of procedures to evolve quantum algorithms

Gepp, Adrian; Stocks, Philip

Published in:
Genetic Programming and Evolvable Machines

DOI:
[10.1007/s10710-009-9080-7](https://doi.org/10.1007/s10710-009-9080-7)

Licence:
Other

[Link to output in Bond University research repository.](#)

Recommended citation(APA):
Gepp, A., & Stocks, P. (2009). A review of procedures to evolve quantum algorithms. *Genetic Programming and Evolvable Machines*, 10(2), 181-228. <https://doi.org/10.1007/s10710-009-9080-7>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

For more information, or if you believe that this document breaches copyright, please contact the Bond University research repository coordinator.

A Review of Procedures to Evolve Quantum Algorithms

Adrian Gepp and Phil Stocks
School of Information Technology
Bond University
Queensland, Australia
{adgepp,pstocks}@bond.edu.au

May 29, 2008

Abstract

There exist quantum algorithms that are more efficient than their classical counterparts; such algorithms were invented by Shor in 1994 and then Grover in 1996. A lack of invention since Grover's algorithm has been commonly attributed to the non-intuitive nature of quantum algorithms to the classically trained person. Thus, the idea of using computers to automatically generate quantum algorithms based on an evolutionary model emerged. A limitation of this approach is that quantum computers do not yet exist and quantum simulation on a classical machine has an exponential order overhead. Nevertheless, early research into evolving quantum algorithms has shown promise.

This paper provides an introduction into quantum and evolutionary algorithms for the computer scientist not familiar with these fields. The exciting field of using evolutionary algorithms to evolve quantum algorithms is then reviewed.

1 Introduction and Overview

Quantum algorithms and evolutionary algorithms are two increasingly popular research fields. Although still relatively new areas of research, there is a large number of publications in both areas. The idea of using evolutionary algorithms to produce quantum algorithms, known as evolving quantum algorithms, has not been pursued to the same extent as research in the two fields independently. This new field of evolving quantum algorithms is in an interesting development stage with the potential to dramatically change the area of quantum computing.

This is a review paper in the field of *evolving quantum algorithms using genetic programming*. The objectives of this paper are to

- provide a detailed introduction to the field of quantum computing,
- introduce the concept of evolutionary algorithms, specifically genetic programming, and
- discuss the application of genetic programming to quantum algorithms, including a comprehensive summary of work in this new field.

These objectives will be addressed in order; this paper is directed at readers with general knowledge of computer science, but not necessarily of quantum computing or evolutionary algorithms.

2 Introduction to Quantum Computing

Since Feynman’s famous speech delivered in December 1959 in which he mentioned the possibility that “sub-microscopic” computers could be built [20], the world has seen great advances in computing power and computer miniaturization, especially in the last two decades¹. Nonetheless, there are limits to computer miniaturisation with current microchip technology [50] and we still, as yet, have not realised Feynman’s sub-microscopic prediction. The desire for “sub-microscopic” computers is fueled by the appeal of the ability to significantly increase the efficiency of storing, copying, transmitting and processing information using a computer that will not occupy significant physical space.

The two major approaches that have been proposed for sub-microscopic computing are *quantum computing* and *DNA computing*, which have been widely discussed theoretically and are the subject of a large number of theoretical and empirical studies. DNA or molecular computing, not to be confused with evolutionary or genetic algorithms, essentially involves using DNA molecules instead of microchips. This allows the information-processing capabilities of organic molecules to replace digital switching primitives and achieve this sub-microscopic size [50]. As this section of this paper only introduces quantum computing, more information on DNA computing is contained in the excellent book by Păun, Rozenberg and Salomaa [51] or the shorter article by Groß [28].

2.1 Brief History of Quantum Computing

Quantum computing is derived from a link between quantum mechanics, computer science and classical information theory [39]. Essentially quantum computing can be viewed as developing algorithms that will run on quantum computers, which are based on quantum mechanics. In the 1920s, the theory of quantum mechanics was proposed: the well-known major contributors were Born, Dirac, Heisenberg, and Schrödinger [39]. However, it wasn’t until 1980 that the area of quantum computing truly began, when Benioff [6] presented his *quantum Turing machine*. It proved that quantum systems could coherently perform computation to solve problems.

Two years later in 1982, Feynman [19] observed that quantum computing could not always be simulated efficiently on a classical computer, which led to speculation that quantum computing could be more efficient than classical computing in some cases. In 1985, Deutsch [14] further developed Benioff’s quantum Turing machine and suggested that quantum computers might be able to solve problems that had no efficient solution on a classical or probabilistic Turing machine. He also presented an example oracle problem², now known as Deutsch’s problem, that was solved more efficiently on a quantum computer than by any classical algorithm. In addition, four years later Deutsch [15] developed another, more popular, model for quantum computing known as the *quantum circuit model*. Nevertheless, neither Deutsch’s problem nor its later developed generalisation Deutsch-Jozsa’s problem [16] proved to be practically useful, so the field continued to develop slowly.

Peter Shor [57] shocked the computing world in 1994 when he presented two efficient quantum algorithms, for factoring integers and the ‘discrete-log’ problem, for which there was, and still is, no efficient classical counterpart. Researchers had been searching for an efficient factoring algorithm for over 20 years, and most people were confident that this algorithm did not exist as the most efficient classical algorithm discovered is the number field sieve [41] that requires exponential time. In contrast, Shor’s factoring algorithm requires only polynomial time. Shor’s algorithm was also the first non-oracle

¹A more detailed and structured account of Feynman’s idea was given in his book [21].

²An oracle problem is such that some properties of a function are trying to be discovered, whereby that function is given as a black box. The code of a black box function is unknown, but the function’s results for given inputs can be calculated.

quantum algorithm produced. Shor's algorithm was inspired by Simon's quadratic-time quantum algorithm [60] that solves an oracle problem, which requires exponential time on a classical computer³. Nevertheless, all the focus was on Shor's factoring algorithm, because of its extremely important applications in cryptography. All classical cryptography techniques, such as RSA, would be easily breakable with a quantum computer running Shor's algorithm. Therefore, it was this algorithm that spurred much interest and research in quantum computing.

Even with this increase in research as a consequence of Shor's breakthrough, the only breakthrough since then has been Grover's search algorithm [29] developed in 1996. Grover's algorithm searches an n -element unstructured list in $O(\sqrt{n})$, compared to the classical $O(n)$. Thus, Grover's algorithm only provides a quadratic speed-up, but it is a major breakthrough due to its wide-spread application in search problems. Since 1996, the only new quantum algorithms have been variants of Shor's and Grover's algorithms. Possible reasons for the slow discovery of quantum algorithms are presented in Section 2.4.3.

2.2 Preliminaries

Depending on the reader's knowledge of linear algebra, computer science and quantum computing, some of the topics introduced in this section may already be familiar to the reader. Therefore, as these topics can be read independently, it is possible to just refer to those topics of personal interest.

2.2.1 Dirac Notation

Quantum states can be described as vectors, which are by convention expressed in the notation invented by Paul Dirac in 1958 [18]. The basic two elements of the Dirac notation are called bras and kets. Standard column vectors in a Hilbert Space⁴ are represented by kets, such as $|v\rangle$ for a column vector v . The matching bra is a row vector, denoted $\langle v|$, which represents the conjugate transpose, or dual, of v . The conjugate transpose of v is defined as the transpose of the vector, in which each element is the complex conjugate of the corresponding element in v . These single vector representations can then be combined to represent such operations as the inner and outer product of vectors; for example, the inner and outer product of v and ψ are written as $\langle v|\psi\rangle$ and $|v\rangle\langle\psi|$ respectively.

2.2.2 Tensors and Tensor Products

In essence, a tensor is a geometrical identity that is the generalisation of n -dimensional vectors with their associated linear operators. Tensors can also be represented as multi-dimensional arrays. However, from a quantum computing standpoint, the tensor product is the most important operation that can be performed on tensors: the tensor product of v and ψ is denoted $v \otimes \psi$.

The tensor product is the most general bilinear operation: it is a generalisation of the matrix product operation, whereby all values contained in an operand are multiplied independently with all values in the other operand. Thus, the result of $v \otimes \psi$ has the following properties:

$$\begin{aligned} \text{Rank}(A \otimes B) &= \text{Rank}(A) + \text{Rank}(B) \\ \text{Dimension}(A \otimes B) &= \text{Dimension}(A) \times \text{Dimension}(B) \end{aligned}$$

³Simon's algorithm, like Deutsch's, lacks practical application.

⁴Chapter 16 of Hardy and Steeb's book [32] contains a good introduction to Hilbert Space.

The actual operation of the \otimes operator is more clearly described by the following example:

$$A \otimes \begin{pmatrix} a & b \\ c & d \end{pmatrix} = \begin{pmatrix} A \times a & A \times b \\ A \times c & A \times d \end{pmatrix}$$

A complete treatment of the mathematics of tensor products is unnecessary here considering the scope of this paper; for a more complete treatment of tensor products refer to Hungerford's book [34]. The application of tensor products in quantum computing is discussed in Section 2.2.4 and 2.2.5.

2.2.3 Quantum Bit

A quantum bit, commonly referred to as a qubit, is the basic unit of information in quantum computing: it is analogous to the bit in classical computation. Similar to a classical bit, a qubit has two computational *basis states*, sometimes referred to as *eigenstates*, usually represented as $|0\rangle$ and $|1\rangle$, that correspond to a bit's 0 and 1 states respectively. However, unlike a classical bit, a qubit can be in a superposition of these two basis states. Due to this superpositional nature of a qubit it can be thought of as a 2-dimensional vector (in complex vector space) of length one with the above two basis states as orthonormal vectors:

$$\begin{aligned} |0\rangle &= \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\ |1\rangle &= \begin{pmatrix} 0 \\ 1 \end{pmatrix} \end{aligned}$$

Thus, a qubit in a superposition of basis states can be described as being a unit vector that lies between the two basis states. A qubit in an arbitrary state is therefore expressed as a linear combination of basis states:

$$\alpha|0\rangle + \beta|1\rangle \text{ or } \begin{pmatrix} \alpha \\ \beta \end{pmatrix},$$

where α, β are complex numbers and $|\alpha|^2 + |\beta|^2 = 1$

Nevertheless, just like a classical bit, a quantum bit can be read to get its value, which can only be one of the basis states. However, the outcome is not deterministic as in classical computing, but rather probabilistic. Given the above expression for a qubit, the probability of the qubit being measured in each basis state is determined by the values of α and β , which are referred to as *amplitudes* or *probability amplitudes*. The actual probability of the qubit being measured in a basis state is the square of the corresponding amplitude; for example, the probability of the qubit being measured as a $|0\rangle$ is $|\alpha|^2$. Hence, the need for the condition $|\alpha|^2 + |\beta|^2 = 1$.

As the two amplitudes of a qubit in a basis state are 1 and 0 (order irrelevant) it will have 100% probability of being measured in the state it is in and 0% probability of being measured in the alternative basis state; that is, a qubit in a basis state is measured deterministically. On the other hand, measurement of a qubit in a superposition of basis states changes the qubit into the basis state in which it is measured. The ramifications of this are that subsequent measurements of a qubit initially in superposition are deterministic and will yield the same outcome with 100% probability. Therefore, a qubit in a basis state can be thought of as a classical resource; in contrast, a qubit in a superposition of basis states can be thought of as a purely quantum resource [74].

2.2.4 Multiple Qubits

Individual bits in classical computation combine through cartesian product, but quantum bits combine through tensor product [54]. Taking the simplest case of multiple

qubits, two qubits, we can find the possible basis states as the tensor product of the two individual sets of basis states:

$$\begin{aligned} & \{|0\rangle, |1\rangle\} \otimes \{|0\rangle, |1\rangle\} \\ &= \{|0\rangle \otimes |0\rangle, |0\rangle \otimes |1\rangle, |1\rangle \otimes |0\rangle, |1\rangle \otimes |1\rangle\} \quad \text{or more concisely:} \\ & \{|00\rangle, |01\rangle, |10\rangle, |11\rangle\} \end{aligned}$$

The final concise version represents an extension of Dirac notation whereby $|q_0q_1 \dots q_n\rangle$ represents the basis for qubits $q_0 \dots q_n$. Thus, with k written as a binary number, a two-qubit quantum system can be described as:

$$\sum_{k=0}^3 \delta_k |k\rangle \text{ or } \begin{pmatrix} \delta_0 \\ \delta_1 \\ \delta_2 \\ \delta_3 \end{pmatrix}, \text{ where } \delta_k \text{ are complex numbers and } \sum_{k=0}^3 |\delta_k|^2 = 1$$

Thus, an n-qubit quantum system can be written as $\sum_{k=0}^{2^n-1} \delta_k |k\rangle$, where $\sum_{k=0}^{2^n-1} |\delta_k|^2 = 1$ and $|\delta_k|^2$ is the probability of the system being measured in the basis state $|k\rangle$. It is possible to measure a sub-set of the qubits in a multi-qubit system. It is also possible to determine such information as whether two qubits are equal without learning their value; however, these more complex measurements are equivalent to a transformation of the quantum system followed by a standard measurement to determine the basis state of each qubit [54], and consequently it is common practise to also refer to standard measurements.

2.2.5 Entangled States

A suprising and non-intuitive aspect of quantum computing is that there are quantum states that can not be described in terms of their individual component qubits. These states are known as *entangled states*, and the individual qubits within entangled states are known as *entangled particles*. The reason for the use of the term ‘entangled’ should become obvious in the next paragraph.

We established in the previous section that single qubits combine using the tensor product operator. Therefore, two qubits in superposition combine as follows:

$$\begin{aligned} |q_1\rangle \otimes |q_2\rangle &= (\alpha_1|0\rangle + \beta_1|1\rangle) \otimes (\alpha_2|0\rangle + \beta_2|1\rangle) \\ &= \alpha_1\alpha_2|00\rangle + \alpha_1\beta_2|01\rangle + \alpha_2\beta_1|10\rangle + \beta_1\beta_2|11\rangle \end{aligned}$$

However, there are entangled states that can not be described by the above formula. The canonical examples are the Bell states:

$$\begin{aligned} |\psi_{00}\rangle &= \frac{|00\rangle + |11\rangle}{\sqrt{2}} & |\psi_{01}\rangle &= \frac{|01\rangle + |10\rangle}{\sqrt{2}} \\ |\psi_{10}\rangle &= \frac{|00\rangle - |11\rangle}{\sqrt{2}} & |\psi_{11}\rangle &= \frac{|01\rangle - |10\rangle}{\sqrt{2}} \end{aligned}$$

In order to write ψ_{00} in terms of the above two-qubit equation $\alpha_1\beta_2$ must equal 0. This subsequently implies that either $\alpha_1\alpha_2 = 0$ or $\beta_1\beta_2 = 0$; however, it is clear that neither of these terms are zero. Thus, ψ_{00} is an entangled state (the other Bell states are shown to be entangled with similar reasoning). An example of how a non-entangled state can be expressed in terms of its component qubits is the state $\frac{1}{\sqrt{2}}(|10\rangle + |11\rangle) = |1\rangle \otimes \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$

Another way of viewing entangled states is that measurement of one qubit has an effect on the other qubits. Take ψ_{00} as an example again, if the first qubit is measured as $|0\rangle$ then we know that the second qubit will also be measured as $|0\rangle$. In contrast, measurement of the first and second qubit in our non-entangled state will always yield $|0\rangle$ for the first and a fifty-fifty chance of either basis state for the second, regardless of the measurement of the measurement order. Some famous research explaining this

observation includes that of Einstein, Podolsky and Rosen: an accurate and concise summary of this work is given by Rieffel and Polak [54]. It is clear that there is no classical counterpart of entangled states. However in quantum computing, it is important to realise that entangled states are treated no differently from any other quantum states. Nevertheless, they are extremely important and some very interesting applications of entangled state have been found such as *dense coding*, *quantum teleportation* (see Section 2.7.1) and secure *quantum key distribution*, which are discussed by Rieffel and Polak [54]. A detailed analysis of entangled states, their importance and a C++ implementation, is provided by Hardy and Steeb [31].

2.3 Quantum Computers

Quantum computers still remain abstract machines despite the large amount of money and time that has been invested into building them. Quantum computers with a few qubits have been developed; however, a scalable general-purpose quantum computer is yet to be built. The largest quantum computer developed thus far consists of seven qubits [69], which was built by a group of scientists from Stanford University. It was constructed using Nuclear Magnetic Resonance (NMR) and tested using Shor's factoring algorithm [75]. In February 2007, the Canadian company D-Wave publicly demonstrated 'Orion', which they claim is a 16-qubit quantum computer. However, details of Orion's inner workings are unknown as no academic papers have been published about Orion⁵. Thus, there is academic speculation about whether Orion is truly a *quantum* computer. There have been many other approaches to building a quantum computer, including optical photon spins, quantum dots, cavity quantum electrodynamics and ion traps. NMR and ion traps are the most advanced approaches to date; however, all of these techniques have had scaling problems. Therefore, it is commonly argued whether or not quantum computers of a significant scale (greater than 100 qubits) will ever be practically realised. Pellizzari [48] presents a general overview of the requirement for building a quantum computer, including its challenges, various approaches and future outlook; however, this article was published in 1998 and consequently does not include the new NMR approach. Hardy and Steeb [32] provide a concise summary of NMR, and Nielsen and Chuang [47] provide a more detailed look at all of the above mentioned approaches.

There are certain characteristics that an operational quantum computer must have regardless of the approach adopted for its construction. Different publications outline various characteristics that are somewhat similar [32, 42, 47]. These characteristics have been combined and summarised below:

Qubit Representation There must be a representation of a qubit (dynamic two-state object) such as the up and down spins of a proton.

State Preparation The system must reliably start in a known initial state. It is sufficient and common to start with all qubits in the same state (conventionally $|0\rangle$); thus, this requirement has not been a major issue.

State Transformation There must be a mechanism to efficiently transform a quantum system into other states following the rules of quantum mechanics. The most common way this has been implemented is by the use of *quantum gates*, which will be discussed in Section 2.6.1.

Measurement There must be an efficient and reliable way to measure qubits in the system.

System Isolation The system must be isolated from the environment to prevent the superposition of states from decaying quickly in the common environment, which is a phenomenon known as decoherence or quantum noise. This is a problem as

⁵One of many news articles about Orion can be viewed at <http://arstechnica.com/articles/paedia/hardware/quantum.ars>.

algorithms may not be able to be applied in time if the states decay too quickly. This issue of decoherence is a major hurdle for the actualisation of a quantum computer [54, 37, 59]. Zurek [84] provides a comprehensive and simple discussion of decoherence.

It is almost inevitable that complete isolation will not occur and there will be some decoherence. As a partial solution, research is being undertaken on quantum error-correcting codes (QECCs). Since the first 9-qubit QECC was presented by Shor in 1995 [58] research and development in this area has been making noticeable progress; the most recent work presents the most time-efficient QECC (although not the most space-efficient) [23]. Despite progress being made in developing quantum error-correcting code, this method will always come at a cost of extra qubits, which consequently increases the number of qubits needed to implement any given algorithm. This is the reason QECC is solution for only a small amount of decoherence. Nielsen and Chuang [47] present a comprehensive section on the area of quantum error-correction, which has been summarised by Landry [37]. Furthermore, Hardy and Steeb [32] discuss the latest algorithms for developing quantum error-correcting code.

Whether a scalable quantum computer will ever be built obviously has huge impacts for the field of quantum computing, as well as classical computing. Nevertheless, quantum computing is currently a popular research field and this paper simply views a quantum computer as an abstract machine. Furthermore, while this paper does not discuss quantum mechanics directly, most of the basic postulates of quantum mechanics will be discussed within Section 2 of this paper. For a structured coverage of the exact postulates of quantum mechanics refer to Nielsen and Chuang's book [47].

2.4 Power of Quantum Computing

After more than 50 years of using the classical physics paradigm (which the classical Turing machine embodies) to build a theory of (standard) computation, quantum physics provides a different paradigm that is arguably more powerful than standard computation [4]. The possibility of harnessing the power of quantum parallelism (discussed in the following section) and the identification of entangled states that are without a classical counterpart (discussed in Section 2.2.5) were the first indications that quantum computing could allow faster information processing than classical computing. Furthermore in 1985, Deutsch [14] proposed that his quantum Turing machine could not only simulate quantum systems better than classical methods, but it may also be able to solve classical problems significantly faster than classical, and possibly probabilistic, Turing machines. This essentially suggested a violation in the strong Church-Turing thesis that any algorithmic process can be simulated on a classical Turing machine (CTM) in at worst a polynomial slowdown.

2.4.1 Quantum Parallelism

Quantum parallelism is the term used to describe the potential parallel computing power of quantum computing. It is intuitive that an increase in time efficiency can be gained by using parallel processors. Furthermore, in classical computation, an exponential time efficiency increase requires an exponential increase in the number of processors or physical space. However, as qubits can represent a superposition of two different states, in quantum computation, a linear increase in physical space generates an exponential increase in parallelism, and consequently an exponential increase in time efficiency. This is what is known as *quantum parallelism* [16]. As shown in Section 2.2.3, a qubit can be represented as a two-dimensional vector. In classical physics, n two-dimensional vector objects form a $2n$ -dimension vector space. However due to quantum parallelism this is not the case in quantum computing in which n qubits form a 2^n -dimension vector

space. This exponential growth implies a possible exponential increase in the information processing speed of quantum computers over classical computers. The concept of how quantum parallelism is applied to create parallel computation is conveniently thought of as shown below in a two-qubit example:

$$\begin{array}{ccc} \begin{pmatrix} \delta_0 \\ \delta_1 \\ \delta_2 \\ \delta_3 \end{pmatrix} & \otimes & T \\ \text{All Inputs} & \otimes & \text{One Transformation} \end{array} = \begin{array}{c} \begin{pmatrix} \delta_0 \otimes T \\ \delta_1 \otimes T \\ \delta_2 \otimes T \\ \delta_3 \otimes T \end{pmatrix} \\ \text{All Outputs} \end{array}$$

A set of input qubits in a superposition of basis states essentially has all the possible inputs encoded in them. Therefore, a *single* transformation can be applied to generate a set of output qubits in a superposition of basis states, which represents all possible outputs. Thus, all output states have been computed and assigned a probability of measurement on a quantum computer in the same time it takes to compute the output for a single input state on a classical computer.

Although all outputs can be produced using only one transformation, we know from Sections 2.2.3 and 2.2.4 that only one output can be extracted upon measurement. Furthermore, the particular output that will be extracted is unknown as measurement is probabilistic. Thus, a quantum algorithm must manipulate quantum parallelism so that the desired result is extracted when measured. This is a difficult and non-intuitive task that no classical programming technique can solve. The two best-known approaches are [54]

- Measure common properties of all the outputs. This approach is used in Shor's factoring algorithm whereby the period of the outputs is measured
- Increase the amplitudes of the basis states of interest to increase the probability that they will be measured. This technique is used in Grover's search algorithm.

In addition, there are other restrictions on quantum computing that are discussed in the next section.

2.4.2 Unitary Restriction

With the exception of measurement, all transformations in a quantum system must be unitary. This is specified in Schrödinger's equation, which governs the dynamics of a quantum system which is not being measured. To be precise, Schrödinger's equation enforces that quantum systems must change state in a way that preserves orthogonality⁶, and the only transformations that maintain orthogonality in a complex vector space, such as quantum systems, are unitary transformations. The transformation matrix M is a unitary matrix if $MM^* = I$, where M^* is the conjugate transpose of M . A unitary transformation can be thought of as a rotation in complex vector space [54]: this makes sense as vector rotations will maintain the angles between all vectors and preserve orthogonality.

An important consequence of quantum transformations being unitary is that they must be reversible, as unitary matrices are invertible. For computing to be reversible, the current state must uniquely determine the previous state, for all computational states [7]. However, this is not a problem as Bennett [7] showed that all classical computation can be transformed into reversible computation without a polynomial expenditure in time or space. Another important consequence of unitary transformations is that their linearity implies that quantum states can not be cloned, which is known as the *no-cloning theorem*. An easily understandable version of the original proof of this theorem by Wootters and Zurek [80] is given by Rieffel and Polak [54]. It is important

⁶To be orthogonal means to be linearly independent: a precise treatment of the orthogonality condition is given by MathWorld [available at <http://mathworld.wolfram.com/OrthogonalityCondition.html>].

however to note that this theorem only applies to qubits in an unknown state: qubits in a basis state or a known state of superposition⁷ can be cloned, but qubits in an unknown state can not be cloned as the information about their amplitudes will be lost.

2.4.3 The Potential

There is a well-established view that quantum computing will only yield an exponential speedup in problems whose structure avoids the need to process exponentially many cases [12]. That is, a brute force approach to NP-complete problems will not succeed with the aid of quantum parallelism; Fortnow and Rogers [24] also firmly question whether quantum computing is more powerful than classical computing, but they still maintain it is a potentially powerful model of computation. However, whether quantum computation can efficiently solve NP-complete problems in polynomial time remains an open question [54]. With the discovery of a few quantum algorithms there is no doubt that certain problems can be solved more efficiently on a quantum computer than on a classical computer. The number and diversity of problems that can be solved more efficiently on a quantum computer is still unknown and the subject of much current research.

In reality, it is not extremely difficult to develop a working quantum algorithm. However, quantum algorithms are only of interest if they are *more* efficient than their classical counterparts. This is one of the main reasons that more quantum algorithms have not been discovered. It is simple enough to simulate a classical algorithm on a quantum computer; however, producing a quantum algorithm that is superior to its classical counterpart requires the use of truly quantum effects⁸, which is extremely complex. Pessimists speculate that the lack of discovery of quantum algorithms is due to a lack of application of quantum computing. However, it seems more likely the reasons are the non-intuitive nature of creating quantum algorithms to the classically trained world and the fact that only superior quantum algorithms are of interest. Nevertheless, it has not been conclusively proven that quantum computation is more powerful than classical computation (this issue will be discussed more in the following section). The potential of quantum computation is however extremely high as there may even be an exponential speed gap between a quantum and classical computer [27], as is the case for Shor's factoring algorithm [57].

2.5 Complexity of Quantum Algorithms

Quantum computing is a unique type of probabilistic computing, which has caused the creation of new complexity classes [54]. The most interesting quantum complexity class is bounded-error quantum polynomial time (*BQP*), which is the quantum equivalent of the classical bounded-error probabilistic polynomial time (*BPP*) complexity class. *BPP* contains all languages for which there is a probabilistic Turing machine such that it gives the right answer with bounded probability, which means greater than or equal to $\frac{2}{3}$ of the time. Note that the probability $\frac{2}{3}$ is simply a norm as it can be replaced with any number between $\frac{1}{2}$ and 1 without altering the complexity class: if the probability is set higher than $\frac{1}{2}$ any previously *BPP* algorithm can simply be rerun to amplify the probability [59]. Using this definition of *BPP* and substituting 'quantum Turing machine' for 'probabilistic Turing machine' we obtain the definition of the complexity class *BQP*.

It is a significant result in quantum complexity theory that Bernstein and Vazirani proved that $BPP \subseteq BQP \subseteq PSPACE$ [9], but the open question of extreme importance in quantum computing is whether *BQP* is strictly larger than *BPP*. It is important to note that proving $BPP \subset BQP$ implies that $BPP \subset PSPACE$, which is

⁷A known state of superposition is where all the amplitudes are known.

⁸The term quantum effects is used as an umbrella term to encompass effects such as superposition and entanglement that exist in quantum computing, but not classical computing.

not currently known. Thus, proving that quantum computers are more powerful than classical computers would also represent a breakthrough in classical computing complexity theory. However, it also indicates that proving $BPP \subset BQP$ is probably extremely difficult. Simon’s quadratic-time quantum algorithm [60] to solve a problem with known classical complexity of exponential-time suggests that BQP is strictly larger than BPP , but it is not a rigorous proof as it only shows $BPP \subset BQP$ for oracle problems [59]. Grover’s search algorithm [29] also puts forward a case that BQP is strictly larger than BPP ; however, it remains an open question [39]. That is, it has not yet been proven that quantum computers have more capabilities than classical computers.

Further details on quantum complexity classes are outside the scope of this paper, but more information on the BQP complexity class is presented by Fortnow and Rogers [24], and [8, 77] detail quantum computing complexity in time and space respectively.

2.6 Models of Quantum Computation

The first quantum computation model was the quantum Turing machine (QTM) introduced by Benioff [6] and Deutsch [14], which can efficiently simulate every classical Turing machine (CTM) with a polynomial-bounded overhead [27]. The QTM is also the commonly used basis for defining universality. In essence, the QTM has three deterministic CTM tapes where the two extra tapes are needed to make the computation reversible (the set of tape states also changes to reflect the nature of quantum computing). The QTM is described in more detail by Gramß [27].

There are also many other mathematically equivalent models of quantum computation, just as there are many different models of classical computation. Moreover, there are other potential models of quantum computation that have not been explored [59]. A quantum state machine [30] has been proposed that is an extension of classical finite state machines in which amplitudes are added to transitions to represent quantum parallelism; a universal quantum cellular automata model [73] is another quantum computing model. Recently, a universal quantum lambda calculus [74] has also been suggested that is based upon the classical lambda calculus. Essentially, a reversible lambda calculus is established with the use of minor reduction steps, and the quantum parallelism is captured as subexpressions representing superposition that can not be reduced/observed (no-cloning theorem).

This section of this paper focuses on the quantum circuit model, due to its popularity and simplicity. The quantum circuit model, also known as the quantum gate array, was introduced by Deutsch in 1989 [15]. The circuit model is mathematically equivalent to the QTM [83], and is the most popular model of quantum computation. The reason for the popularity difference is that the QTM suffers from the same complexities as the CTM; it does not satisfy a simple algebra and it can be cumbersome to use as it requires ‘word at a time’ thinking while keeping track of control variables such as tape states [74]. In contrast, quantum circuits are easier to understand as they are graphical diagrams, which can also be manipulated algebraically. Furthermore, other models such as the quantum lambda calculus are too new to be widely popular. The quantum circuit model will be explained in more detail in the following sub-section.

2.6.1 Quantum Circuit Model and Quantum Gates

Similar to classical circuits, the quantum circuit model consists of wires, and gates that act upon wires. In the quantum circuit model, each wire represents a qubit⁹ and each individual gate has the same number of input qubits as it has output qubits (due to the reversibility of quantum computing). Thus, an advantage of this model is that the diagram itself can be essentially thought of classically, with the data on the wires representing the quantum nature of the computing. As the number of qubits (wires) is required to be constant, an important restriction of this model is that a quantum circuit

⁹By convention a qubit begins in a basis state, usually assumed to be $|0\rangle$.

can only compute a function with a specific length domain. Therefore, for functions with arbitrary length inputs a *family* of quantum circuits is required; that is, a quantum circuit for each input length is required.

The complexity of a quantum circuit depends on the total number of gates and qubits (wires). However, as gates are commonly restricted to act upon one, two or three qubits¹⁰, the number of gates in a quantum circuit is both a reasonable and the usual measure of complexity. This restriction does not affect the universality of quantum circuits; for example, the set of all one-qubit gates and the CNOT gate (explained later in this section) are universal for quantum computing. Further note that an interesting three-qubit gate is the Toffoli gate that is in fact universal for quantum computing as presented by Deutsch [15]. In addition, it has more recently been shown that a large number of two-bit quantum gates are universal [3]. An excellent review of quantum gates and the respective universality issue is contained in [4].

Quantum gates are in fact unitary transformation matrices (see Section 2.4.2), and a matrix that acts upon n -qubits will be a $2^n \times 2^n$ square matrix due to quantum parallelism. Quantum parallelism (see Section 2.4.1) is the reason for a linear increase in the number of qubits causing an exponential increase in the size of a quantum system. To familiarise the reader with both the gates and the quantum circuit model¹¹, some important elementary gates are detailed below using the quantum circuit model.


Three important single-qubit gates are the identity (I), negation (X), and Hadamard (H) transformations, which are detailed below¹²:

$$\begin{aligned}
 I: & \quad \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} & |0\rangle \rightarrow |0\rangle; |1\rangle \rightarrow |1\rangle \\
 X: & \quad \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} & |0\rangle \rightarrow |1\rangle; |1\rangle \rightarrow |0\rangle \\
 H: & \quad \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} & |0\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle); |1\rangle \rightarrow \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)
 \end{aligned}$$

It is important to notice that the Hadamard gate creates an equal superposition of states, and when applied to n -qubits it creates a superposition of all 2^n states (this is the reason for its prevalence in quantum algorithms). The workings of these gates can be easily verified with linear algebra; for example, the negation transformation on $|0\rangle$:

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \times 0 + 0 \times 1 \\ 1 \times 1 + 0 \times 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

These single-qubit transformations are represented in the quantum circuit model by an appropriately labeled box as shown below:



The controlled-not (CNOT) is an extremely important two-qubit gate. Its importance comes from the fact that the results of a CNOT transformation can not be expressed as the tensor product of two qubits; thus, a CNOT gate can be used to create


¹⁰There is some inconsistency in the literature about the maximum number of qubits that a gate should act upon; for example, Shor [59] states the maximum should be two qubits while Spector et al. [63] state that the maximum is a few qubits.

¹¹Variations in the quantum circuit notation of different gates exist, but they are only minor and should be understood by a reader who has read this paper.


¹²Due to linearity, transformations are fully specified by their effect on the basis states [54].

entangled states. The workings of the CNOT gate are relatively simple: it swaps the target qubit (represented by a cross) if the control qubit (represented by an open circle) is $|1\rangle$. The mapping and quantum circuit notation for the CNOT gate are as shown.

$$\begin{aligned} |00\rangle &\rightarrow |00\rangle & |01\rangle &\rightarrow |01\rangle \\ |10\rangle &\rightarrow |11\rangle & |11\rangle &\rightarrow |10\rangle \end{aligned}$$



The universal 3-qubit Toffoli gate is essentially a controlled-controlled-not gate, where the target qubit is only flipped if the two control qubits are $|1\rangle$; It is represented in the circuit model as follows.



2.7 Example Problems and Algorithms

The algorithms that will be presented in this section are quantum teleportation, and an overview of Grover's algorithm. These algorithms have been chosen for their ease of understanding and their ability to convey important quantum algorithm concepts. Pittenger's book [49] contains excellent, succinct treatments of the following quantum algorithms: Deutsch-Jozsa's, Simon's, Grover's, Shor's and the Quantum Fourier Transformation.

This section will also include definitions of Deutsch's and the scaling majority-on problems. Quantum algorithms to solve these problems have been evolved by Spector et al. using genetic programming (see Section 4.3.3).

2.7.1 Quantum Teleportation Algorithm

Quantum teleportation¹³, uses two classical bits to transfer one quantum bit of information ($|q1\rangle$) from A(lice) to B(ob) without ever being anywhere in between¹⁴. This does not violate the no-cloning theorem as the unknown qubit is teleported not copied; that is, after the quantum teleportation the qubit no longer exists with Alice. The quantum circuit model for teleportation is as shown below in Figure 1.

The first step creates Bell's first entangled state $|\psi_{00}\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}$ (discussed in Section 2.2.5) by putting the first qubit into equal superposition using the Hadamard gate and then modifying the second according to the value of the first using the CNOT gate. Now, given that $|q1\rangle$ is unknown, it can be represented as $\alpha|0\rangle + \beta|1\rangle$, and the quantum state after the first two transformations is

$$\frac{1}{\sqrt{2}} (\alpha|0\rangle (|00\rangle + |11\rangle) + \beta|1\rangle (|00\rangle + |11\rangle))$$

¹³This explanation of the quantum teleportation algorithm is based on the explanation given by Landry [37].

¹⁴Quantum teleportation of one qubit has been realised experimentally [11].




Figure 1: Quantum Teleportation Algorithm

Alice and Bob are each given one of these two entangled particles and the final two transformations are performed as shown, and the next CNOT is applied to make the quantum state

$$\frac{1}{\sqrt{2}} (\alpha|0\rangle (|00\rangle + |11\rangle) + \beta|1\rangle (|10\rangle + |01\rangle))$$

Thus, the information about $|q1\rangle$ is now contained in the entangled pair, and the final Hadamard transformation is performed to yield

$$\begin{aligned} & \frac{1}{2} (\alpha (|0\rangle + |1\rangle) (|00\rangle + |11\rangle) + \beta (|0\rangle - |1\rangle) (|10\rangle + |01\rangle)) \\ &= \frac{1}{2} [|00\rangle (\alpha|0\rangle + \beta|1\rangle) + |01\rangle (\alpha|1\rangle + \beta|0\rangle) + |10\rangle (\alpha|0\rangle - \beta|1\rangle) + |11\rangle (\alpha|1\rangle - \beta|0\rangle)] \end{aligned}$$

Alice then measures the two qubits she has and sends the result encoded in two classical bits $\{00,01,10,11\}$ to Bob. Using the previous equation, Bob knows the state of his qubit in terms of $|q1\rangle$. Therefore, he can apply a simple one-qubit gate (if needed) to convert his qubit to $|q1\rangle$ as defined below (using the previous equation):

Alice's Result	Bob's Qubit	Transformation Required
$ 00\rangle$	$\alpha 0\rangle + \beta 1\rangle$	<i>none</i>
$ 01\rangle$	$\alpha 1\rangle + \beta 0\rangle$	<i>negation(X)</i>
$ 10\rangle$	$\alpha 0\rangle - \beta 1\rangle$	$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$
$ 11\rangle$	$\alpha 1\rangle - \beta 0\rangle$	$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$

For example, in the case of $|01\rangle$, after transformation Bob's qubit equals

$$\begin{aligned} X (\alpha|1\rangle + \beta|0\rangle) &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} b \\ a \end{pmatrix} \\ &= \alpha|0\rangle + \beta|1\rangle \\ &\equiv |q1\rangle \end{aligned}$$

Thus, an unknown qubit in a superposition of states can be transported from A to B with only two classical bits of information.

2.7.2 Grover's Algorithm

Grover's algorithm searches an unstructured n -element list in $O(\sqrt{n})$ time compared to its classical counterpart which is $O(n)$. The first step in Grover's algorithm is to apply a Hadamard gate to n -qubits initially in state $|0\rangle$. The Grover operator (or Grover iteration) is then applied multiple times ($O(\sqrt{n})$) to these n -qubits as well as extra qubits that are required as workspace (referred to as the *oracle workspace*). Essentially, the Grover operator increases the amplitude of the basis states that are being searched for and decreases the amplitudes of the other states. This is done by rotating the current




Figure 2: The result from applying the Grover operator: relabelled from Figure 6.3 in Nielsen and Chuang’s book [47]. The three numbers in boxes indicate the initial (1), intermediate (2), and final (3) state vectors when the Grover operator is applied. $|\beta\rangle$ is the superposition of all solutions to the search problem, and $|\alpha\rangle$ is a state orthogonal to $|\beta\rangle$. All vectors in this diagram are unit vectors, but $|\alpha\rangle$ and $|\beta\rangle$ are shown to be longer to improve clarity.

state vector (of the n-qubits) towards the superposition of all solutions to the search problem ($|\beta\rangle$) as shown in Figure 2.

The Grover operator consists of

1. Applying the oracle search function (using the oracle workspace) to the n-qubits.
2. Applying Hadamard transformations to all n-qubits
3. Performing a conditional phase shift
4. Applying Hadamard transformations to all n-qubits

The initial oracle application reflects the current state of the n-qubits about $|\alpha\rangle$ to move it from its initial state 1 to state 2, as shown in Figure 2. The remaining three operations (Hadamard, conditional phase shift, Hadamard) then reflect the state 2 vector about state 1. Overall, this achieves a rotation of θ towards $|\beta\rangle$ (state 3). Repeating the Grover operator $O(\sqrt{n})$ times rotates the state vector very close to $|\beta\rangle$, the superposition of all solutions. Hence, when the state is measured with reference to the computational basis it reveals a solution to the search problem with extremely high probability¹⁵. This is the essence of Grover’s algorithm; an excellent detailed analysis of Grover’s algorithm and its recent variations is given in chapter 6 of Nielsen and Chuang’s book [47].

2.7.3 Deutsch’s and Scaling Majority-On Problems

Deutsch’s problem is determining whether a given oracle, or black box, function is uniform or balanced, given that the oracle must be either one of these. The uniform property requires an oracle to always return 0 or always return 1; the balanced property requires an oracle to return an equal number of 0s and 1s over all possible inputs. The scalable majority-on problem is an extension of Deutsch’s problem where the oracle is an arbitrary boolean function and the problem is to determine whether the majority of the outputs are 1. Quantum algorithms to solve these problems have been evolved by Spector et al. using genetic programming (see Section 4.3.3).

¹⁵The exact probability of observing a solution depends upon the number of solutions in the search space.

2.8 Further Reading

This section presents the major summary references in the field of quantum computing that complement the specific references given in previous subsections. However, most of the references given here have been previously referenced in this paper.

There are currently a large, and ever increasing, number of publications on quantum computing, and consequently there is a lot of duplication and work of varying standards. For issues on quantum computing, Nielsen and Chuang's book [47] appears to be the most comprehensive and well-structured publication, which is considered to be the most significant reference. A very good introductory reference is provided by Landry [37], who summarised the major areas in Nielsen and Chuang's book. Another excellent and recommended review reference is Rieffel and Polak's work [54], which contains a more comprehensive and detailed review of quantum computing than Landry's paper [37]. Furthermore, in contrast to [37, 54], Gramß [27] provides a summary of quantum computing from a quantum Turing machine perspective, rather than the quantum circuit model. However, many (although not all) of the topics included in [27, 37, 54] have been covered in this paper.

There are two excellent quantum computing books that also contain detailed information on how to actually simulate quantum algorithms on a classical computer: Williams and Clearwater's book [79] comes with Mathematica¹⁶ notebooks that simulate well-known quantum algorithms such as Shor's factoring algorithm, and Hardy and Steeb's book [32] contains Java and C++ code for some simple quantum simulations such as generating entangled qubits.

The majority of the articles referenced in this quantum computing section are freely available at the Los Alamos preprint server: <http://xxx.lanl.gov/archive/quant-ph>. This site also provides an excellent place to search for old and new articles within the field of quantum computing.

3 Introduction to Evolutionary Algorithms

Optimisation problems can be characterised by two sets of parameters: feedback parameters to optimise according to a target solution, and free parameters to modify in order to approach the desired solution [10]. The optimisation algorithm alters the free parameters while controlling the quality of the solution by the feedback parameters; different optimisation techniques perform this search in different ways. When traditional optimisation techniques are used to search vast, complex and unknown spaces there are extremely complex constraints and multimodal problems¹⁷ [32]. Thus, traditional optimisation techniques are not well-suited for these types of problems, so alternative approaches have been researched. Evolutionary Algorithms (EAs) are one of the alternative methods that have gained significant popularity as general-purpose optimization and search tools [32]. EAs are probabilistic search algorithms that are heavily based upon Darwinian evolution, as a proxy for the process of species evolution in nature. The central concept of Darwinian evolution is that individuals in a population have heritable characteristics that influence their probability of producing offspring, that is, future generations. EAs have extended this theory slightly to state that characteristics of 'better' individuals will increase their likelihood to produce offspring, which is a variation on Darwin's 'survival of the fittest' principle [13]. The idea behind this is an attempt to converge to the 'best' individual, which essentially is the paradigm of search and optimisation.

EAs were initially used as optimisation tools for engineering problems, and were developed independently by several computer science researchers in the 1950s and 1960s [25]. Since then, the number of applications for EAs has become diverse and has grown

¹⁶Mathematica is a comprehensive mathematical software package, details are available on their website: <http://www.wolfram.com/products/mathematica/index.html>.

¹⁷Multimodal problems are problems that arise in cases with a large number of locally optimum solutions.

at a fast rate [10]; for example, there are EA applications in financial forecasting, predicting protein structure, predicting the primeness of numbers and in developing computer programs. In contrast to the application aspect of EAs, despite much theoretical research there has been modest progress in EA theory over the last 20 years compared with that of neural networks, another biologically motivated form of computation [5]. Nevertheless, various types of EAs have been developed and the types presented in the following subsection now form the backbone of the EA field [46]. Regardless of what type of EA is used the basic elements are almost identical [2]; these common elements are [39]

- Populations of individuals representing solutions to the problem at-hand, which allows parallel searching.
- Ways to manipulate solutions, which can be either
 - Mutation (inspired by the biological process of the same name) operators, which implement innovative change
 - Recombination/Crossover (inspired by the biological process of gamete production and sexual reproduction) operators, which implement conservation of characteristics through rearrangement
- A measure for determining the quality of a solution, usually referred to as a fitness function
- A method of selection that uses the fitness function to select individuals for the next generation

3.1 Types of EAs




Figure 3: Relationships between different types of Evolutionary Algorithms

EAs and Genetic Algorithms (GAs) are terms that are sometimes incorrectly used interchangeably. However, EA is an umbrella term that includes all algorithms that incorporate the idea of Darwinian evolution, and GAs are just one type of EA. The different types of EAs are shown in Figure 3. GAs, invented by Holland in the 1960s [33], are the most prominent EAs [32], with the other major types of EAs being Evolutionary Programming (EP) and Evolution Strategies (ES). Essentially, EP and ES operate on and change the phenotype (observable properties) of individuals, while GAs operate on the genotype (genetic construction) of individuals [22]. To further explain the concept of phenotype, a change in phenotype can be viewed as a change in the behavior, physiology or morphology without altering the genetics. In addition, a further difference between ES and GAs is that ES places an emphasis on mutation [10], whereas GAs place a higher weighting on recombination. However, having stated the differences between these types of EA, it is important to note that these distinctions are not strict, as overlap does occur [10]. For example, GAs have been applied to real numbered genomes [70], which overlaps into the ES field. For a more thorough review of EP and ES, which also includes a discussion on GAs, please refer to Bäck’s book [1].

In 1992, Koza introduced a new type of EA known as Genetic Programming (GP) [36]; Bornholdt [10] suggests that the reason GP arrived late into the EA field may have been due to the need for greater computing power. GP is a type of GA whereby

the search space is reduced to solely include computer programs. GP techniques have been valuable in evolving structures other than computer programs, but the fact that individuals in a GP model are computer programs is the most defining feature of GP [61]. Nevertheless, GP has extended the idea of genotype manipulation from GAs to include variable length chromosomes; that is, the representation of individuals in a population can be of varying size.

3.2 GA and GP Algorithm Structure

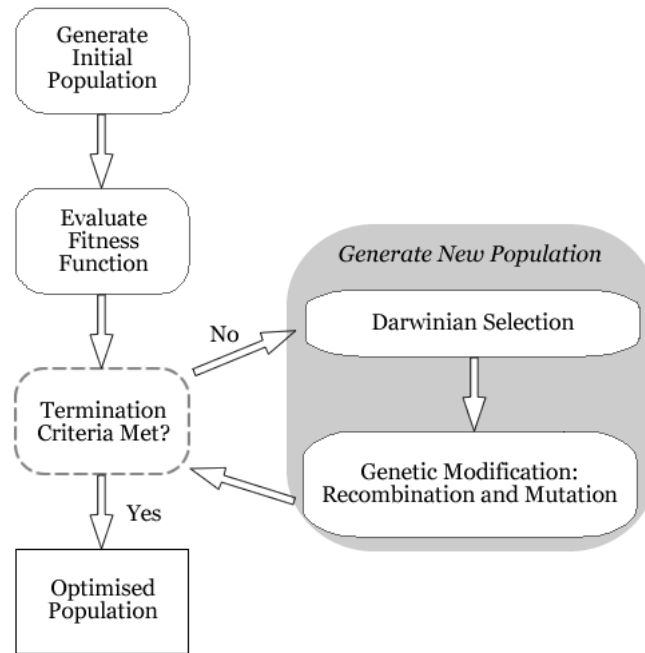


Figure 4: A general flow diagram of GAs and GP algorithms.

The overall structure of GAs and GP algorithms (GPAs) is shown in Figure 4. Initially, a starting population is generated; this is usually done randomly so that it results in a diverse population that simulates all naturally occurring populations. However, before an initial population can be created there must be a way to represent individuals. Traditionally, as supported by Holland [33], individuals, or solutions, were encoded as binary strings; modern GAs (and consequently GPAs) use tailored encoding to suit the problem being solved [45]. The number of individuals in the initial population is another parameter that must be set according to the specific problem at hand.

After an initial population is created the evolutionary loop begins. The evolutionary loop consists of the following repeating sequence:

1. Evaluate the fitness of each individual in the population. The fitness, or quality, of each individual (x) is measured by a predetermined fitness function ($f(x)$) with a real number codomain, that is, it returns a real number.
2. Select individuals based on fitness levels to be parents of the next generation, as in Darwinian Selection. The three main selection methods used in GAs and GPAs are
 - (a) *Fitness-Proportional* selection, where individuals are selected using a *roulette wheel* concept in which each individual (i) has a sector of size proportional to its fitness (f_i). Assuming that a lower fitness function corresponds to a superior fitness, this means each individual has a probability of being selected

equal to

$$\frac{1 - f_i}{\sum_i (1 - f_i)}$$

- (b) *Ranking* selection, which simply involves selecting the best x (predefined parameter) number of individuals, according to a predefined function based on their fitness.
 - (c) *Tournament* selection, which involves randomly selecting a number (predefined parameter) of individuals to compete in the tournament. The losers of the tournament, based on fitness, are then excluded from selection, and all remaining individuals are selected. This is however, a much weaker form of selection than the ranking selection.
3. Generate the next population through genetic modification. The mutation operator is applied to a single individual which results in a random change to its representation (genome) at one or more positions; it is applied with a certain low probability, set as a parameter. The crossover operator is applied to two individuals (parents) and produces a new individual (child) for the next generation, which is a combination of the representations (genomes) of the original individuals. The crossover operator is also applied with a probability set as a parameter, but this probability is usually much higher than the mutation probability. Thus, cross-over can obviously be performed in many ways and depends on the representations of the individuals. Two possible approaches are (1) using the beginning of the representation from one parent and the end from the other parent, or (2) reusing one parent's representation with a middle section exchanged for a section of the other parent. One common classification of different crossover operators can be separated into the following two categories:
- *Fixed-length* crossover, where the children have the same length representations as their parents. Assuming that individuals are a sequence of instructions, an example of fixed-length crossover is taking the first 4 instructions from parent A (length 6) and the last 2 instructions from parent B (length 6), which would produce a child that had 6 instructions, the same as its parents.
 - *Variable-length* crossover, which simply means that the result of the crossover operation may result in a different length child. For example, taking the first 4 instructions from parent A (length 6) and the last 5 instructions from parent B (length 6) produces a child that had 9 rather than 6 instructions.

This loop terminates when sufficiently optimised individuals¹⁸ have been evolved or other predetermined termination criteria are met, such as the maximum number of generations have been evolved.

3.2.1 Further Reading

For further information on GAs, Bornholdt [10] presents a succinct review, and Hardy and Steeb [32] provide a good review of the implementation of GAs, which includes many Java and C++ code examples. In addition, a very comprehensive analysis of GAs from a mathematical standpoint is given by Vose [76]. More detail on GP in general is covered in Koza's books [36, 35].

4 Evolving Quantum Algorithms

Genetic Programming (GP) is defined as developing algorithms based on Darwinian evolution to automatically generate computer programs. Thus, as quantum computing is a formal computational model, GP can be used to evolve quantum computing programs (represented as quantum algorithms); this field will be hereafter referred to

¹⁸The desired level of optimisation is set as a parameter in terms of the fitness function.

as ‘Evolving Quantum Algorithms’. Unlike other applications of GP, there has been relatively little work done in the field of evolving quantum algorithms [55].

As was stated in Section 2.4.3, it seems that the major reasons for a lack of discovery of quantum algorithms are that they are difficult and non-intuitive to generate, and that only faster than classical quantum algorithms are of interest. Therefore, it would seem appropriate that we look to computers to search for quantum algorithms. However, as the search space of quantum algorithms is vast, complex and reasonably uncharted, using traditional search techniques is problematic (see Section 3). Hence the motivation to use GP to evolve quantum algorithms, with the hope that GP’s power to randomly search vast, complex and unknown spaces can discover many superior quantum algorithms.

4.1 Simulation Limitations

In order to evolve quantum algorithms the fitness of different individuals (quantum algorithms) must be assessed; the fitness of a solution is assessed by running the algorithm. However, as quantum computers have not yet been built (Section 2.3), quantum algorithms must be simulated on a classical machine. However, due to quantum parallelism simulating quantum algorithms on a classical computer comes at a potentially exponential simulation cost [63, 55, 40]. Therefore, as inefficiencies in simulation will only multiply rapidly during the evolutionary process, an efficient quantum computer simulator is a necessity in this field [39]. Spector [61] introduces a tested quantum computer simulator termed QGAME, which is based upon the quantum circuit model detailed in Section 2.6.1; Sabry [56] outlines an interesting quantum simulation model in the functional language Haskell, and there are also many quantum computer simulators that vary in quality available on the Internet. An example code fragment of the quantum simulator used by Spector et al. [63] explains a possible representation of gates in the programming language LISP, which is demonstrated below:

```
(HADAMARD 0) ;; apply Hadamard transformation to qubit 0  
(CNOT 1 2) ;; apply CNOT gate with control (qubit 1) and target (qubit 2)
```

In addition, recall that Hardy and Steeb [31] provide some C++ and Java code for simulating different quantum effects. Nevertheless, due to the simulating inefficiency only small quantum algorithms can be simulated and evolved. In this sense small quantum algorithms means that the following must be restricted in size: number of qubits, number of gates (commonly referred to as the length of a quantum circuit), and the number of iterations of the algorithm (with varying input) needed to determine the fitness of the algorithm. Hence, discovering small, yet scalable, quantum algorithms is the ultimate result from evolving quantum algorithms.

4.2 Program Structures

Within GP, the individuals (computer programs) must have a defined structure so that genetic operators can be applied to them. In most applications of GP, the computer programs consist of basic building blocks, referred to as primitives or genomes, such as constants, operators, problem-specific functions and inputs [39]. When evolving quantum algorithms the main components of the quantum algorithm are quantum gates. These primitives can then combine to form complete computer programs in different structures; the types of program structures that have been applied in the field of evolving quantum algorithms are *tree*, *linear*, and the hybrid *linear-tree*. The choice of an appropriate program structure appears to be a key ingredient to the success of the GP evolution [39].

The original GP program structure, as outlined by Koza [36, 35], was based on the standard tree structure. Quantum gates and their arguments are represented as parent nodes and their children respectively. Children nodes can be quantum gates

themselves¹⁹, and leaf nodes can be constants, zero-argument functions, qubit indices or other program inputs. This structure can be easily translated into a functional programming language where the functions are the parent nodes and the function’s arguments are its children. The program can also be executed with a post-order traversal (left-subtree, right sub-tree, root), which means no further memory is required as function arguments are always locally accessible. It is also important to note that it is trivial to convert a tree-structured quantum program into the quantum circuit model, where the order of gates in the circuit model will be defined by the post-order traversal²⁰.

Although Koza’s tree structure provides elegant adaptations to different program sizes and shapes [38] and allows for easy expression in any functional programming language, there is no guarantee it is the most appropriate representation for all problems [66]. An alternative to the tree structure, is a linear structure that unsurprisingly consists of a linear sequence of instructions. Spector et al. [66] outline reasons that suggest tree-based program structures offer no advantages over linear program structures, but have additional complexity. They state that quantum algorithms are linear, so a linear program structure appears to be more appropriate than a tree structure. Operators or functions in the linear based model then get their arguments from external memory. The linear-tree structure is a combination of these two structures, where there is a linear structure and branching function at each node of the tree structure. During execution of the tree structure, the linear structure of each node visited is executed. The tree structure execution begins by visiting the root node and executes only one path to a leaf node: the next node visited is determined by the branching function at the current node. In quantum computing, partial measurements of the entire quantum state are appropriate branching functions [39]; therefore, the next node depends on the result of the partial measurement. An example branching function is if the measurement of qubit n yields $|0\rangle$ or $|1\rangle$, then the left or right subtree will be visited next respectively.

It should also be noted that there is a *graph*, and consequently a *linear-graph*, GP program structure that has been applied to areas other than quantum algorithms. There are many types of graph structures: the well-known PADO graph structure is presented by Teller and Veloso [71]. It seems likely that the reason this structure has not been used to evolve quantum algorithms is that the complexity cost is greater than that of other structures, but this approach does not seem to represent quantum effects better than other structures.

4.3 Previous Studies

In addition to the papers cited, Leier’s PhD thesis [39] is a major source for this section and provides more details on the specific parameters of each study than are given in this section. A new paper by Giraldi, Portugal and Thess [25] also contains, among other things, reviews of Rubinstein’s [55] and Yabuki and Iba’s [81] work; however, Leier’s thesis [39] appears to be the best reference.

4.3.1 Williams and Gray

The book by Williams and Gray [78] is the pioneering work that suggests the use of genetic programming to evolve quantum algorithms. However, the goal of their study was to search more efficiently for alternative quantum circuits to a known quantum algorithm than conducting an exhaustive search. The difference from subsequent studies is that, in this case the overall unitary gate that represents the quantum algorithm was already known. This is still useful as it could be used to search for (hopefully faster) alternative quantum algorithms that solve the same problem [55]. In this case, the

¹⁹The return value of a quantum gate node is a valid index of a qubit it acted upon. For example, the CNOT gate is usually defined to return the index of the control qubit.

²⁰There is also a trivial mapping of a quantum circuit to a tree-structured program, based on establishing a sequence of gates, that is further described in Leier’s thesis [39].

fitness function would include a penalty for circuits that were similar to the known quantum algorithm.

Using a population size of 100, the GP algorithm implemented in this study successfully found a quantum teleportation algorithm (see Section 2.7.1 for details on quantum teleportation), where the send portion was as efficient, and the receive portion more efficient, than the known algorithm; the GP algorithm used:

- A linear program structure
- An approximate universal-gate set consisting of the CNOT gate and two other one-qubit gates
- A three-tuple representation: {quantum gate, [parameters²¹]²², set of qubits acted upon}
- A fitness function based on the evolved circuits similarity to the known solution
- A ranking selection scheme
- Mutation and Crossover operators that act on quantum gate(s)

4.3.2 Yabuki and Iba

Yabuki and Iba [81] developed a GA model²³ that was specifically tailored to evolve the quantum teleportation algorithm as did Williams and Gray [78]. Using a larger population size of 5000, the GA algorithm implemented in this study successfully found a quantum teleportation algorithm (see Section 2.7.1 for details on quantum teleportation) that has at least 3 fewer gates than any non-evolved quantum teleportation algorithm; the GA algorithm used:

- A linear program structure
- The same gate set as Williams and Gray [78] (chosen as the problem was the same)
- A unique *fixed-length* string representation that was specifically tailored to the production of two entangled particles, which is essential in quantum teleportation
- A problem specific fitness function based on the error of the evolved algorithm's output
- A fitness-proportional selection scheme
- Mutation and fixed-length crossover operators

Yabuki and Iba, like Williams and Gray, have approached evolving quantum algorithms from an optimisation, rather than a search, standpoint. They have focused on using a tailored evolutionary algorithm to find a more efficient quantum algorithm for a given existing quantum algorithm. Furthermore, along with Williams and Gray's [78] results, it has been shown that there is potential in this area, particularly for the quantum teleportation algorithm.

4.3.3 Spector et al.

Spector has recently published the first book [62] about using GP to automatically evolve quantum algorithms, which makes reference to all of his previous work with a more detailed introduction into the field. Spector et al. [63, 64, 65, 66] have conducted extensive research on evolving quantum algorithms. Three different GP models were outlined and applied to various problems in [66]. The common elements of all three models are the

- Standardised fitness function (detailed in [63]) that takes into account three components:

²¹Only discrete parameters were allowed.

²²[] indicates optionality.

²³It is technically not a GP model due to the fixed length representation of individuals.

1. *Hits*: total number of fitness cases minus the number of fitness cases for which the quantum circuit yields the correct answer with less than 0.48 probability
2. *Correctness*: $\frac{\sum_{i=1}^{numCases} \max(0, error_i - 0.48)}{\max(hits, 1)}$. This formula deliberately ignores any circuits with errors less than 0.48, so that the focus is on producing correct quantum algorithms, not on improving the probability of correctness of already correct quantum algorithms²⁴
3. *Efficiency*: number of quantum gates / 100,000, which makes sure that this component will always be less than 1.

Spector et al. recommend that these components be combined lexicographically, whereby quantum circuits are only compared on Correctness if their Hits are the same, and similarly only compared on Efficiency if the Hits and Correctness are the same.

- Tournament selection method with a tournament size of five individuals

The first GP model defined in [66] uses a standard tree program structure, and is in fact a summary of the results presented in Spector et al.’s initial paper [63]. This GP model was applied to two oracle problems, namely Deutsch’s problem and the scaling majority-on problem (defined in Section 2.7.3). Using a population size of 10,000, this tree structure GP algorithm was used to evolve a better-than-classical quantum algorithm for Deutsch’s problem, but did not evolve a quantum algorithm to solve the majority-on problem better than any classical algorithm. Along with the common elements shown above, this tree structure GP algorithm also used

- A set of gate building functions as well as iteration structures and arithmetic operators
- Functions in prefix notation, analogous to functional languages such as LISP in which the quantum algorithms were simulated
- Mutation operators, and two types of crossover operators: a variable length operator design for the tree structure and a reproduction operator to produce a child exactly the same as a parent.

The other two GP models use two types of linear program structures. As stated in Section 4.2, a linear structure stores arguments in external memory. The first type of linear structure, called the stack-based linear GP, uses a global stack for temporary data storage. This GP model successfully evolved a quantum algorithm to solve the four-item database search problem faster than any classical algorithm; moreover, the evolved quantum algorithm is as efficient as, and almost identical to, Grover’s search algorithm. Besides the program structure, the differences between this GP model and the tree structure model discussed above are that the

- Various crossover operators are designed to operate on linear, rather than tree, structures; they are also fixed-length operators
- Iteration structures are stack-based, rather than tree-based
- Gate building functions do not return values on to the stack
- Arithmetic functions take their arguments from, and return their result to, the stack
- A *no-op* operator is part of the function set

The primary role of including iteration structures is to produce scalable quantum algorithms, however, some non-scalable quantum algorithms are of interest [66]. The second type of linear GP model is tailored to evolving these non-scalable quantum algorithms as there is no iteration structure. Furthermore, Spector et al. [66] suggest that there is no major reason for quantum gates to share parameter values, which

²⁴Here Spector et al. are using the term correct to mean correct greater than or equal to 48% of the time.

consequently means there is no reason for data storage. Thus, in contrast to the stack-based linear structure, the second type of linear structure implemented, called the stack-less linear GP, only contains the gate building functions and has no external memory for temporary data storage. This GP structure was used to evolve a quantum algorithm to solve the And-Or Query problem faster than any previously known classical or quantum algorithm. The And-Or Query problem is to determine whether the boolean function $(f(0) \vee f(1)) \wedge (f(2) \vee f(3))$ is true or false for a two-bit boolean function f .

Spector et al. [64, 65] then presented a modification of their stack-less linear structure GP. The changes to the GP model include

- Using a steady-state GP. All previous GP models applied to quantum algorithms have been generational GP models; the difference is simply that steady-state GP models do not have clearly defined generations. The remainder of the evolutionary process can be considered to be very similar
- Supporting the use of variable-length representations and variable-length crossover operators
- A four component fitness function, in which the components are again combined lexicographically; the components are the
 1. *Misses* component, which is equivalent to the *Hits* component used in [66]
 2. *Expected-Queries* component that considers the number of oracle calls (defined in [65]), which is a tailored modification for evolving quantum algorithms to solve oracle problems
 3. *Max-Error* component, which is the maximum probability of getting a wrong answer in any fitness case, and is similar to the *Correctness* component used in [66]
 4. *Num-Gates* component, which is equivalent to the *Efficiency* component used in [66]
- Including a one-qubit measurement gate so that partial measurements can be made, which are key to several known quantum algorithms such as Shor's factoring algorithm
- The added ability to distribute the evolutionary process across multiple workstations to decrease execution time

This modified GP model was again applied to the And-Or Query problem, with better results than the initial stack-less linear structure model. From these improving results, Spector et al. deduced that they were successfully improving their GP model to evolve quantum algorithms, and that the stack-less linear structure is probably the best structure that has been developed for evolving quantum algorithms. Furthermore, the successful research by Spector et al. summarised above has shown that there definitely is a degree of potential in evolving quantum algorithms.

4.3.4 Rubinstein

Using a population size of 5000, Rubinstein [55] used his generational GP algorithm to successfully discover the most efficient known quantum algorithms to produce two to five maximally entangled qubits²⁵, that is, qubits of the form $\frac{1}{\sqrt{2}}(|0\dots 0\rangle + |1\dots 1\rangle)$. The GP algorithm used

- A linear program structure (with no external storage)
- An unspecified set of quantum gates that include CNOT, Hadamard and importantly an 'Observe' gate that can measure one or many qubits, which is known to be a vital technique in several known quantum algorithms, such as Shor's factoring algorithm

²⁵Two maximally entangled qubits is in fact the first Bell state as described in Section 2.2.5.

- A modification of the three-tuple representation, where the quantum gate and its parameters and qubit operands are encoded into a bit string, which is the standard representation for an individual in GPs
- A fitness function based solely on the error of the evolved quantum system. The fitness was calculated using the following formula, where i is an input case of k total cases, j is a basis state in a quantum system of n qubits, and o and d are the observed and desired amplitudes respectively.

$$error = \sum_i^{k-1} \sum_j^{2n-1} (o_{ij} - d_{ij})$$

The error obtained from this function is then divided by the highest error of any individual evolved to obtain a standardised fitness that lies in the range between 0 (optimal) and 1.

- A fitness-proportional selection scheme
- Mutation operators (with low probability), and crossover operators that act upon all parts of the bit string representation, that is, gates, parameters and qubits. Gate cross-over is variable-length, but parameter and qubit crossover are fixed-length, as most quantum gates act upon a fixed number of qubits with a fixed number of parameters

From the quantum algorithms produced for the maximum entanglement problem with two to five qubits, deductions were made about an arbitrary sized maximum entanglement production circuit. Thus, this study shows that there is potential for GP to produce small, yet scalable, quantum circuits that can be converted into large scale quantum circuits.

4.3.5 Lukac et al.

Lukac and Perkowski [43] present a general GA for evolving quantum circuits. Their algorithm used:

- An encoding system where quantum circuits are represented as an array of strings of quantum gates. Each element of the array represented a specific point in time (after the previous element and before the next element) in the quantum algorithm when gates could act upon qubits. The string of ordered gates at each element corresponded to the order of qubits, such that the the first gate acted upon the first qubit and so on. This system had no extra parameters to identify which qubits a gate acted upon, thus it was only possible with the introduction of a no operation gate. Their encoding system also deliberately allowed for parallel evaluation of individuals to potentially decrease fitness evaluation time.
- A large gate set comprising various one, two and three qubit gates, including a one-qubit ‘wire’ gate that performs ‘no operation’. The reason for the diverse gate set is Lukac and Perkowski’s focus on evolving arbitrary quantum algorithms.
- A fitness function considering the correctness of the quantum algorithm. This fitness function is similar to those in previous papers [55, 78], but in this case a fitness level of 1 corresponds to maximum (not minimum) fitness.
- Roulette wheel and stochastic universal sampling selection schemes. Stochastic universal sampling is less biased towards selecting ‘fit’ individuals, thus a more random selection scheme (which is detailed more in [26]).
- Mutation and crossover operators.
- Population sizes of 50 and 100 individuals. The population was deliberately kept small to avoid long fitness evaluation times.

Lucak and Perkowski used their algorithm to evolve various quantum circuits, averaging each experiment over 20 runs. First, they evolved single gate circuits (whereby the target gate was included in the gate set for the experiment) to test their algorithms convergence. All gates were successfully evolved, but for larger number of input-qubits (gate size) longer evolution time was required. Interestingly, mutation with a probability greater than 0.4 was found to decrease both real-time evolution and the number of generations needed to successfully evolve the target gate. Their second experiment was to evolve three quantum circuits consisting of more than one gate: namely, the quantum teleportation algorithm previously evolved by Williams and Gray [78] and three and four maximally entangled qubits as previously evolved by Rubinstein [55]. In all three cases, the desired quantum circuit was evolved in similar or less time than previously published. It was also noted that while the number of generations required to evolve quantum circuits increases exponentially as the number of qubits increase, the real-time evolution increases at a much slower rate. Higher probability of mutation was again found to decrease the number of generations and real-time for successful evolution.

While not discovering any new or further optimised quantum algorithms, Lucak and Perkowski have established benchmarks for the evolution of various small quantum algorithms, ranging from 1 to 4 qubits. Furthermore, for the three composite quantum algorithms (> 1 qubit) evolved, their GA has performed equally or more efficiently (in terms of time of evolution) than previous studies. Lukac et al. [44] have furthered their research by investigating implementing their GA on a quantum computer. The specific encoding used in this GA lends itself to be computer in parallel, thus it should be implementable on quantum technologies, such as nuclear magnetic resonance (NMR).

4.3.6 Leier

Leier, in his recent PhD thesis [39], presented two GP models for evolving quantum algorithms, one linear and one linear-tree structure. Both models are very similar to the successful stack-less linear structure GP model developed by Spector et al. [65, 64]. The major difference is the four component fitness function of Leier’s model, in which the components are combined through different weightings, rather than lexicographically. The actual components are also slightly different, but have all been previously been mentioned by Spector et al.; they are: *misses*, *max-error*, *correctness* and *num-gates*. Further differences in the linear-tree model, caused by its structural difference, are that branching functions are partial measurements and the crossover operators that act on the linear and tree substructures are included. The inherent inclusion of partial measurements and the added flexibility of the linear-tree structure was the motivation for creating a GP model that was not strictly linear.

Both GP models were applied to the Deutsch-Jozsa (D-J) and 1-SAT problems²⁶. The Deutsch-Jozsa problem is essentially a scalable version of Deutsch’s problem; 1-SAT also has a known better-than-classical quantum algorithm solution (Hogg’s algorithm). The linear-tree GP was able to find a quantum algorithm essentially the same as the known algorithms for both the 1-SAT and D-J problem, although some evolutionary runs did not produce a solution to the D-J problem. The interesting finding from this study was that intermediate partial measurements had no noticeable positive effect; similarly, the added flexibility of the linear-tree structure did not add a benefit over the strict linear structure which also evolved solutions to both problems. However, it is probable that these findings can not be generalised past these small problems that have relatively simple quantum solutions. Nevertheless, this study was the first to show that linear-tree structure GP models can be used to evolve quantum circuits.

Leier also made some other interesting observations; the most important of these are listed below:

²⁶Note that the result of the linear-tree models applied to 1-SAT is also contained in a paper by Leier and Banzhaf [40].

- Using the linear structured GP model, Leier showed that there was an increase in the efficiency of the evolutionary process when using a pre-evolved, rather than a random, initial population for both the D-J and 1-SAT problem. This pre-evolution involved feeding evolved individuals from a smaller problem instance into the initial population.
- According to Leier’s research crossover is not as important as commonly thought. This finding conflicts with the traditional GP approach, where crossover is performed with much greater probability than mutation with the idea of multiplying and distributing better solutions over the population.
- Even though Leier made a point of emphasising that his GP models were not designed to produce scalable quantum algorithms, scalability was achieved: the algorithms evolved using Leier’s models, which apply to n -qubits, can be easily manually scaled to apply to $(n+1)$ -qubits. However, this is in fact not surprising as the known quantum algorithms for both application problems are scalable.

4.3.7 Ding et al.

Ding et al. [17] recently presented a new framework for evolving quantum circuits that is designed for both quantum algorithm discovery and optimisation. This framework uses a Hybrid Quantum-Inspired Evolutionary Algorithm (HQEA), which was motivated by GP and detailed in Yang’s Masters thesis [82]. Ding et al.’s approach used:

- A fixed length numerical encoding of quantum circuits, compared with encoding in symbols as done in previous works. Encoding the quantum circuits with numerical values was necessary to take advantage of the HQEA algorithm.
- A gate set comprising the Hadamard, CNOT, Phase and $\pi/8$ gates²⁷. Their approach is not limited to these gates, but gates in their quantum algorithms are confined to one-qubit and adjacent two-qubit gates. For example, the control qubit of the CNOT gate must be adjacent to the qubit undergoing the NOT operation.
- A fitness function that considers both the cost and correctness of quantum algorithms (as also done in Reid’s Masters thesis [52]). The fitness function used, where lower fitness is better, was:

$$fitness = reward \times (actualcost - satcost) + punish \times (1 - correctness)$$

The *satcost* represents a satisfactory algorithms cost, whereby if it is set high or low then the evolution is inclined towards discovery or optimisation respectively. Algorithm cost was calculated with one-qubit gates, two-qubit gates, and the wire costing 1, 2 and 0 respectively. However, using this same framework algorithm cost could be more accurately computed in terms of the monetary costs using different quantum technologies.

- A fitness-proportional selection scheme
- Mutation and crossover operators.

Ding et al. tested their approach on evolving 2 and 3 entangled qubits, as well as the controlled-phase gate, which confirmed that lower *satcosts* resulted in more optimised quantum algorithms. However, more generations of evolution were required to evolve optimised algorithms. Further experiments were conducted on evolving 2 entangled qubits as research into the appropriate values for the reward-punish factors. This revealed that a large punish:reward ratio (≥ 5 for *satcost* = 6) is required to evolve correct quantum algorithms, and larger punish is required for larger *satcost*. In addition, Ding et al. present a faster method for matrix multiplication with tensor product, to increase the efficiency of evaluating individuals. This faster method can

²⁷Within the scope of this paper, Phase and $\pi/8$ gates can be thought of as reflection (about a basis state) and rotation gates respectively.

also be adapted to other evolutionary algorithms. Overall, this new approach has shown promise in evolving and optimising small quantum algorithms, but it has not yet been used to discover a previously undiscovered quantum algorithm.

4.4 Further Reading

The papers mentioned in the previous section provide further reading: Spector et al.'s work that was published in a book [66] is recommended as the first paper to read due to its understandability. Spector's book [62] and Leier's excellent PhD thesis [39] are recommended as the main points of call on the topic of evolving quantum algorithms. Both of these references also contain an introduction to both the field of quantum computing and evolutionary algorithms, which is similar to, but more detailed than, the one given in this paper. Spector's paper [61] is also recommended as an excellent short paper that extends the idea of evolving quantum algorithms using GP to evolving arbitrary computational processes.

4.4.1 Related Applications of GP

Genetic programming has also been applied to some other closely-related applications. Quantum program evolution based on density matrices [68] is a new area that is showing promise. Spector et al. [63] also suggested that there was potential in researching how GP algorithms would be executed on quantum computers, and whether a significant speed-up on fitness evaluation is possible using quantum parallelism. This idea has been reviewed by Giraldi, Portugal and Thess [25] in 2004, and Udrescu et al. [72] have presented new research into implementing GA algorithms on quantum computers. Another of Spector et al.'s suggestions is to evolve hybrid quantum-classical algorithms, whereby the quantum algorithm has classical pre and post data processors [63]. Another somewhat related paper [53] looked into a case of using GAs to evolve the hardware of quantum computers: a set of pulse sequences (which can be thought of as rotations) is evolved for a given quantum logic gate, implemented by Nuclear Magnetic Resonance (NMR). The significance of this is that the shorter and more robust the pulse sequence the more efficient the implementation of a quantum gate or algorithm; furthermore, the evolved set of pulse sequences was superior to any previously known set.

5 Conclusion

Genetic algorithms and programming have been successfully used to analyse and optimise known quantum algorithms. Previous studies have also shown that genetic programming can evolve new quantum algorithms, albeit only small quantum algorithms. Another possible outcome of evolving quantum algorithms is that a new useful idea, such as a meaningful sequence of gates, which will change the way future quantum algorithms are developed manually, may be discovered; however as yet, this has not occurred. Nevertheless, manual quantum algorithm generation has had more success than evolving quantum algorithms [39], although no further breakthroughs have been made manually since Shor's and Grover's algorithms. Furthermore, the fact that quantum computers do not yet exist is a huge limitation on the research field of evolving quantum algorithms, as quantum simulation on a classical machine has an exponential order overhead. This means that only small quantum algorithms can be evolved, which are few in number and generally have little practical application. Thus, Leier and Banzhaf [39, 40] speculate that this field will have a brighter future when quantum computers exist. However, there are still ideas that have not been tested as the field of evolving quantum algorithms is relatively new and has not yet been comprehensively researched [39].

References

- [1] BÄCK, T. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford Univeristy Press, 1996.
- [2] BANZHAF, W., NORDIN, P., KELLER, R., AND FRANCONI, F. *Genetic Programming -An Introduction*. dpunkt Heidelberg and Morgan Kaufmann Publishers, San Francisco, 1998.
- [3] BARENCO, A. A universal two-bit gate for quantum computation. In *Proceedings of the Royal Society of London A 449* (1995), pp. 679–683.
- [4] BARENCO, A., BENNETT, C. H., CLEVE, R., DIVINCENZO, D. P., MARGOLUS, N., SHOR, P., SLEATOR, T., SMOLIN, J., AND WEINFURTER, H. Elementary gates for quantum computation. *Physical Review A 52*, 5 (1995), 3457–3467.
- [5] BELEW, R. K., AND VOSE, M. D., Eds. *Foundations of Genetic Algorithms 4* (1997), Morgan Kaufmann.
- [6] BENIOFF, P. The computer as a physical system: A microscopic quantum mechanical hamiltonian model of computers as represented by turing machines. *Journal of Statistical Physics 22*, 5 (1980), 563–591.
- [7] BENNETT, C. H. Logical reversibility of computation. *IBM Journal of Research and Development 17* (1973), 525–532.
- [8] BENNETT, C. H., BERNSTEIN, E., BRASSARD, G., AND VAZIRANI, U. V. Strengths and weaknesses of quantum computing. *Society for Industrial and Applied Mathematics Journal on Computing 26* (1994), 1510–1523.
- [9] BERNSTEIN, E., AND VAZIRANI, U. V. Quantum complexity theory. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computation* (1993), pp. 11–20.
- [10] BORNHOLDT, S. Genetic algorithms. In *Non-Standard Computation*, T. Gramß, S. Bornholdt, M. Groß, M. Mitchell, and T. Pellizzari, Eds. WILEY-VCH, Weinheim, Germany, 1998, pp. 141–178.
- [11] BOUWMEESTER, D., PAN, J.-W., MATTLE, K., EIBI, M., WEINFURTER, H., AND ZEILINGER, A. Experimental quantum teleportation. *Nature 390* (1997), 575–579.
- [12] BRAUNSTEIN, S. L. Quantum computation: a tutorial. Tech. rep., Department of Computer Science, York University, 1995.
- [13] DARWIN, C. *On the origin of species by means of natural selection or the preservation of favoured races in the struggle for life*. Murray, London, 1859.
- [14] DEUTSCH, D. Quantum theory, the church-turing principle and the universal quantum computer 400. In *Proceedings of the Royal Society of London A* (1985), pp. 97–117.
- [15] DEUTSCH, D. Quantum computational networks. In *Proceedings of the Royal Society of London A 425* (1989), pp. 73–90.
- [16] DEUTSCH, D., AND JOZSA, R. Rapid solution of problems by quantum computation. In *Proceedings of the royal society of london series A* (1992), vol. A439, pp. 553–558.
- [17] DING, S., JIN, Z., AND YANG, Q. Evolving quantum oracles with hybrid quantum-inspired evolutionary algorithm. *ArXiv Quantum Physics e-prints* (October 2006).
- [18] DIRAC, P. *The Principles of Quantum Mechanics*, fourth ed. Oxford University Press, 1958.
- [19] FEYNMAN, R. Simulating physics with computers. *International Journal of Theoretical Physics 21* (1982), 467–488.

- [20] FEYNMAN, R. P. There's plenty of room at the bottom: An invitation to enter a new field of physics. Speech at the annual meeting of the *American Physical Society*, December 1959. It is available online at <http://www.zyvex.com/nanotech/feynman.html>.
- [21] FEYNMAN, R. P., AND GILBERT, D. H., Eds. *Miniaturization*. Reinhold, New York, 1961, pp. 282–295.
- [22] FOGEL, D. B. Phenotypes, genotypes, and operators in evolutionary computation. In *Computational Intelligence: Theory and Applications, 5th Fuzzy Days* (Berlin, 1995), Springer-Verlag, pp. 337–342.
- [23] FORNEY, JR, G. D., AND GUHA, S. Simple rate-1/3 convolutional and tail-biting quantum error-correcting codes. submitted to 2005 IEEE International Symposium on Information Theory, 2005.
- [24] FORTNOW, L., AND ROGERS, J. Complexity limitations on quantum computation. *Journal of Computer and System Sciences* 59, 2 (1999), 240–252. Special issue for selected papers from the 13th IEEE Conference on Computational Complexity.
- [25] GIRALDI, G. A., PORTUGAL, R., AND TRESS, R. N. Genetic algorithms and quantum computation. *CoRR cs.NE/0403003* (2004).
- [26] GOLDBERG, D. E. *Genetic Algorithms in Search, Optimisation, and Machine Learning*. Addison Wesley, 1989.
- [27] GRAMSS, T. The theory of quantum computation: An introduction. In *Non-Standard Computation*, T. Gramß, S. Bornholdt, M. Groß, M. Mitchell, and T. Pellizzari, Eds. WILEY-VCH, Weinheim, Germany, 1998, pp. 141–178.
- [28] GROSS, M. Molecular computing. In *Non-Standard Computation*, T. Gramß, S. Bornholdt, M. Groß, M. Mitchell, and T. Pellizzari, Eds. WILEY-VCH, Weinheim, Germany, 1998, pp. 15–58.
- [29] GROVER, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of the 18th annual ACM symposium on the history of computing* (Philadelphia, Pennsylvania, May 1996), pp. 212–219.
- [30] GUDDER, S. Quantum automata: An overview. *International Journal of Theoretical Physics* 28, 9 (1999), 2261–2282.
- [31] HARDY, Y., AND STEEB, W.-H. Entangled quantum states and a C++ implementation. *International Journal of Modern Physics C* 11 (2000), 69–77.
- [32] HARDY, Y., AND STEEB, W.-H. *Classical and Quantum Computing, with C++ and Java Simulations*. Birkhäuser Verlag, Berlin, Germany, 2001.
- [33] HOLLAND, J. H. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, 1975.
- [34] HUNGERFORD, T. A. *Algebra*. Springer Verlag, New York, 1974.
- [35] KOZA, J. R. *Genetic Programming II*. MIT Press, Cambridge, 1992.
- [36] KOZA, J. R. *Genetic programming: on the programming of computers by means of natural selection*. MIT Press, Cambridge, 1992.
- [37] LANDRY, O. Introduction to quantum computing. From Physics Department of McGill University, April 2004.
- [38] LANGDON, W. B., SOULE, T., POLI, R., AND FOSTER, J. A. The evolution of size and shape. In Spector et al. [67], pp. 163–190.
- [39] LEIER, A. *Evolution of quantum algorithms using genetic programming*. PhD thesis, University of Dortmund: Department of Computer Science, 2004.
- [40] LEIER, A., AND BANZHAF, W. Evolving Hogg's quantum algorithm using linear-tree GP. In *Genetic and Evolutionary Computation – GECCO-2003* (Chicago, 12-16 July 2003), E. Cantú-Paz, J. A. Foster, K. Deb, D. Davis, R. Roy, U.-M.

- O'Reilly, H.-G. Beyer, R. Standish, G. Kendall, S. Wilson, M. Harman, J. Wegener, D. Dasgupta, M. A. Potter, A. C. Schultz, K. Dowsland, N. Jonoska, and J. Miller, Eds., vol. 2723 of *LNCS*, Springer-Verlag, pp. 390–400.
- [41] LENSTRA, A. K., AND LENSTRA, JR, H. W. *The Development of the Number Field Sieve*. Lecture Notes in Mathematics 1554. Springer Verlag, 1993.
- [42] LOSS, D., AND DIVINCENZO, D. Quantum computation with quantum dots. *Physical Review A* 57 (1998), 120–126.
- [43] LUKAC, M., AND PERKOWSKI, M. Evolving quantum circuits using genetic algorithms. In *Proceedings of the 2002 NASA/DOD Conference on Evolvable Hardware* (2002), pp. 177–185.
- [44] LUKAC, M., PERKOWSKI, M., KERNTOPF, P., PIVTORAIKO, M., FOLGHERAITER, M., LEE, D., KIM, H., HWUANGBO, W., WOOK KIM, J., AND CHOI, Y. W. A hierarchical approach to computer aided design of quantum circuits. In *Proceedings of the 6th International Symposium on Representations and Methodology of Future Computing Technology* (2003), pp. 201–209.
- [45] MICHALEWICZ, Z. *Genetic algorithms + Data structures = Evolution Programs*, 3rd ed. Springer - Verlag, New York, 1996.
- [46] MITCHELL, M. *An introduction to genetic algorithms*. MIT Press, 1996.
- [47] NIELSEN, M. A., AND CHUANG, I. L. *Quantum computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [48] PELLIZZARI, T. Quantum computers: first steps towards a realization. In *Non-Standard Computation*, T. Gramß, S. Bornholdt, M. Groß, M. Mitchell, and T. Pellizzari, Eds. WILEY-VCH, Weinheim, Germany, 1998, pp. 141–178.
- [49] PITTINGER, A. O. *An Introduction to Quantum Computing Algorithms*. Birkhäuser, Boston, 2000.
- [50] PÄUN, G., ROZENBERG, G., AND SALOMAA, A. *DNA Computing: New Computing Paradigms*. EATCS Series. Springer-Verlag, 1998, pp. 1–6,43–74.
- [51] PÄUN, G., ROZENBERG, G., AND SALOMAA, A. *DNA Computing: New Computing Paradigms*. EATCS Series. Springer-Verlag, 1998.
- [52] REID, T. On the evolutionary design of quantum circuits. Master's thesis, Waterloo, Ontario, Canada, 2005.
- [53] RETHINAM, M. J., JAVALI, A. K., BEHRMAN, E. C., STECK, J. E., AND SKINNER, S. R. A genetic algorithm for finding pulse sequences for nmr quantum computing. *submitted to Physical Review A* (April 2004).
- [54] RIEFFEL, E., AND POLAK, W. An introduction to quantum computing for non-physicists. *ACM Computing Surveys* (2000).
- [55] RUBINSTEIN, B. I. P. Evolving quantum circuits using genetic programming. In *Proceedings of the 2001 IEEE Congress on Evolutionary Computation (CEC2001)* (2001), IEEE Press, pp. 114–121.
- [56] SABRY, A. Modeling quantum computing in Haskell. In *Haskell '03: Proceedings of the ACM SIGPLAN workshop on Haskell* (New York, NY, USA, 2003), ACM Press, pp. 39–49.
- [57] SHOR, P. W. Algorithms for quantum computation: Discrete log and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science* (November 1994), Institute of Electrical and Electronic Engineers Computer Society Press, pp. 124–134.
- [58] SHOR, P. W. Scheme for reducing decoherence in quantum computer memory. *Physical Review A* 52 (1995), 2493–2496.
- [59] SHOR, P. W. Introduction to quantum algorithms. Notes for talk given for the short course at the January 2000 American Math Society meeting, 2000.

- [60] SIMON, D. On the power of quantum computation. In *Proceedings of the 35th annual IEEE symposium on the foundations of computer science (FOCS)* (Santa Fee, USA, November 1994), IEEE Computer Society Press, pp. 116–123.
- [61] SPECTOR, L. The evolution of arbitrary computational processes. *IEEE Intelligent Systems* (May/June 2000), 80–83.
- [62] SPECTOR, L. *Automatic Quantum Computer Programming: A Genetic Programming Approach*. Genetic Programming Series. Kluwer Academic Publishers, 2004.
- [63] SPECTOR, L., BARNUM, H., AND BERNSTEIN, H. J. Genetic programming for quantum computers. In *Genetic Programming 1998: Proceedings of the Thurd Annual Conference (1998)*, J. R. Koza, W. Banzhaf, K. Chellapilla, K. Deb, M. Dorigo, D. B. Fogel, M. H. Garzon, D. E. Goldberg, H. Iba, and R. L. Riolo, Eds., pp. 365–374.
- [64] SPECTOR, L., BARNUM, H., AND BERNSTEIN, H. J. Quantum circuits for or and and of ors. *Journal of Physics A: Mathematical and General* 33 (2000), 8047–8057.
- [65] SPECTOR, L., BARNUM, H., BERNSTEIN, H. J., AND SWAMY, N. Finding a better-than-classical quantum AND/OR algorithm using genetic programming. In *Proceedings of the 1999 Congress on Evolutionary Computation (1999)*.
- [66] SPECTOR, L., BARNUM, H., BERNSTEIN, H. J., AND SWAMY, N. Quantum computing applications of genetic programming. In Spector et al. [67], pp. 135–160.
- [67] SPECTOR, L., LANGDON, W. B., O-REILLY, U.-M., AND ANGELINE, P. J., Eds. *Advances in Genetic Programming: Volume 3*. The MIT Press, 1999.
- [68] STADELHOFER, R. Solving the parity problem on a mixed state quantum computer. Tech. rep., University of Dortmund, 2004. Available upon request.
- [69] STEPNEY, S., AND CLARK, J. A. Evolving quantum programs and protocols. In *Handbook of Theoretical and Computational Nanotechnology*, M. Rieth and W. Schommers, Eds. American Scientific Publishers, 2005. Overview available at <http://www-users.cs.york.ac.uk/%7Eesusan/bib/ss/nonstd/gqprev05.htm>.
- [70] SURRY, P., AND RADCLIFFE, N. A formalism for real-parameter evolutionary algorithms and directed recombination. In Belew and Vose [5].
- [71] TELLER, A., AND VELOSO, M. PADO: A new learning architecture for object recognition. In *Symbolic Visual Learning*, K. Ikeuchi and M. Veloso, Eds. Oxford University Press, 1996, pp. 81–116.
- [72] UDRESCU, M., PRODAN, L., AND VLĂDUȚIU, M. Implementing quantum genetic algorithms: a solution based on grover’s algorithm. In *Proceedings of the Third Conference on Computing frontiers* (New York, 2006), ACM Press, pp. 71–82.
- [73] VAN DAM, W. A universal quantum cellular automaton. In *Proceedings of PhysComp96* (1996), New England Complex Systems Institute, pp. 323–331.
- [74] VAN TONDER, A. A lambda calculus for quantum computation. *SIAM Journal on Computing* 33, 5 (2004), 1109–1135.
- [75] VANDERSYPEN, L. M. K., STEFFEN, M., BREYTA, G., YANNONI, C. S., SHERWOOD, M. H., AND CHUANG, I. L. Experimental realization of shor’s quantum factoring algorithm using nuclear magnetic resonance. *Nature* 414 (December 2001), 883–887.
- [76] VOSE, M. D. *The simple genetic algorithm: foundations and theory*. The MIT Press, 1999.
- [77] WATROUS, J. Relationships between quantum and classical space-bounded complexity classes. In *13th Annual IEEE Conference on Computational Complexity* (June 1998), pp. 210–227.

- [78] WILLIAMS, C., AND GRAY, A. *Automated Design of Quantum Circuits*. Springer, New York, 1997.
- [79] WILLIAMS, C. P., AND CLEARWATER, S. H. *Explorations in Quantum Computing*. Springer Verlag, 1998.
- [80] WOOTERS, W. K., AND ZUREK, W. H. A single quantum cannot be cloned. *Nature* 299 (1982), 802–803.
- [81] YABUKI, T., AND IBA, H. Genetic algorithms for quantum circuit design, evolving a simpler teleportation circuit. In *GECCO-00: Proceedings of the Genetic and Evolutionary Computation Conference* (San Francisco, July 2000), Morgan Kaufman Publishers, pp. 421–425.
- [82] YANG, Q. The research of a hybrid quantum-inspired evolutionary algorithm. Master’s thesis, Wuhan University, China, 2006.
- [83] YAO, A. Quantum circuit complexity. In *Proceedings of the 34th Annual Symposium on the Foundations of Computer Science* (Los Alamitos, USA, 1993), IEEE Computer Society Press, pp. 352–361.
- [84] ZUREK, W. H. Decoherence and the transition from quantum to classical. *Physics Today* 44 (1991), 36–44.