Characterizations of 1-Way Quantum Finite Automata

Alex Brodsky
Department of Computer Science
University of British Columbia
abrodsky@@cs.ubc.ca

Nicholas Pippenger
Department of Computer Science
University of British Columbia
nicholas@@cs.ubc.ca

February 20, 2008

Abstract

The 2-way quantum finite automaton introduced by Kondacs and Watrous[KW97] can accept non-regular languages with bounded error in polynomial time. If we restrict the head of the automaton to moving classically and to moving only in one direction, the acceptance power of this 1-way quantum finite automaton is reduced to a proper subset of the regular languages.

In this paper we study two different models of 1-way quantum finite automata. The first model, termed measure-once quantum finite automata, was introduced by Moore and Crutchfield[MC00], and the second model, termed measure-many quantum finite automata, was introduced by Kondacs and Watrous[KW97].

We characterize the measure-once model when it is restricted to accepting with bounded error and show that, without that restriction, it can solve the word problem over the free group. We also show that it can be simulated by a probabilistic finite automaton and describe an algorithm that determines if two measure-once automata are equivalent.

We prove several closure properties of the classes of languages accepted by measure-many automata, including inverse homomorphisms, and provide a new necessary condition for a language to be accepted by the measure-many model with bounded error. Finally, we show that piecewise testable languages can be accepted with bounded error by a measure-many quantum finite automaton, in the process introducing new construction techniques for quantum automata.

1 Introduction

In 1997 Kondacs and Watrous[KW97] showed that a 2-way quantum finite automaton (2QFA) could accept the language $L = a^n b^n$ in linear time with bounded error. The ability of the reading head to be in a superposition of locations rather than in a single location at any time during the computation gives the 2QFA its power. Even if we restrict the head of a 2-way quantum finite automaton from moving left, we can still construct a 2QFA that can accept the language $L' = \{x \in \{a,b\}^* \mid |x|_a = |x|_b\}$ in linear time with bounded error. However, if we restrict the head of a 2QFA to moving

right on each transition, we get the 1-way quantum finite automaton of Kondacs and Watrous[KW97], which, when accepting with bounded error, can only accept a proper subset of the regular languages.

If the reading head is classical then quantum mechanical evolution hinders language acceptance; restricting the set of languages accepted by 1-way quantum finite automata with bounded error to a proper subset of the regular languages[KW97].

During its computation, a 1-way QFA performs measurements on its configuration. Since the acceptance capability of a 1-way QFA depends on the measurements that the QFA may perform during the computation, we investigate two models of 1-way QFAs that differ only in the type of measurement that they perform during the computation.

The first model, termed measure-once quantum finite automata (MO-QFAs), is similar to the one introduced by Moore and Crutchfield[MC00]. The second model, termed measure-many quantum finite automata (MM-QFAs), is similar to the one introduced by Kondacs and Watrous[KW97], and is more complex than the MO-QFA. The main difference between the two models is that a measure-once automaton performs one measurement at the end of its computation, while a measure-many automaton performs a measurement after every transition. This makes the measure-many model more powerful than the measure-once model, where the power of a model refers to the acceptance capability of the corresponding automata.

First, we present results dealing with MO-QFAs. We show that the class of languages accepted by MO-QFAs with bounded error is exactly the class of group languages. Consequently, this class of languages accepted by MO-QFAs is closed under inverse homomorphisms, word quotients, and boolean operations. We show that MO-QFAs that do not accept with bounded error can accept non-regular languages and, in particular, can solve the word problem over the free group. We also describe an algorithm that determines if two MO-QFAs are equivalent and prove that probabilistic finite automata (PFAs) can simulate MO-QFAs.

Second, we shift our focus to MM-QFAs. We show that the classes of languages accepted by these automata are closed under complement, inverse homomorphisms, and word quotients. We prove by example that the class of languages accepted by MM-QFAs with bounded error is not closed under homomorphisms, and prove a necessary condition for membership within this class. We also relate the sufficiency of this condition to the question of whether the class is closed under boolean operations. Finally, we show, by construction, that MM-QFAs can accept piecewise testable languages with bounded error and introduce novel concepts for constructing MM-QFAs.

The rest of the paper is organized in the following way: Section 2 contains the definitions of the quantum automata and background information, Section 3 discusses measure-once quantum finite automata, Section 4 discusses measure-many quantum finite automata, and Section 5 summarizes.

2 Definitions and Background

2.1 Definition of MO-QFA

A measure-once quantum finite automaton is defined by a 5-tuple

$$M = (Q, \Sigma, \delta, q_0, F)$$

where Q is a finite set of states, Σ is a finite input alphabet with an end-marker symbol \$, δ is the transition function

$$\delta : Q \times \Sigma \times Q \to \mathbb{C}$$

that represents the probability density amplitude that flows from state q to state q' upon reading symbol σ, the state q_0 is the initial configuration of the system, and F is the set of accepting states. For all states $q_1, q_2 \in Q$ and symbols $\sigma \in \Sigma$ the function δ must be unitary, thus satisfying the condition

$$\sum_{q' \in Q} \overline{\delta(q_1, \sigma, q')} \delta(q_2, \sigma, q') = \begin{cases} 1 & q_1 = q_2 \\ 0 & q_1 \neq q_2 \end{cases} . \tag{1}$$

We assume that all input is terminated by the end-marker \$; this is the last symbol read before the computation terminates. At the end of a computation M measures its configuration; if it is in an accepting state then it accepts, otherwise it rejects. This definition is equivalent to that of the QFA defined by Moore and Crutchfield[MC00].

The configuration of M is a linear superposition of states and is represented by an n-dimensional complex unit vector, where $n = |Q|$. This vector is denoted by

$$|\Psi\rangle = \sum_{i=1}^{n} \alpha_i |q_i\rangle$$

where $\{|q_i\rangle\}$ is the set orthonormal basis vectors corresponding to the states of M. The coefficient α_i is the probability density amplitude of M being in state q_i. Since $|\Psi\rangle$ is a unit vector, it follows that $\sum_{i=1}^{n} |\alpha_i|^2 = 1$.

The transition function δ is represented by a set of unitary matrices $\{U_\sigma\}_{\sigma \in \Sigma}$ where U_σ represents the unitary transitions of M upon reading symbol σ. If M is in configuration $|\Psi\rangle$ and reads symbol σ then the new configuration of M is denoted by

$$|\Psi'\rangle = U_\sigma |\Psi\rangle = \sum_{q_i, q_j \in Q} \alpha_i \delta(q_i, \sigma, q_j) |q_j\rangle.$$

Measurement is represented by a diagonal zero-one projection matrix P where $P_{ii} = [q_i \in F]$. The probability of M accepting string x is defined by

$$p_M(x) = \langle \Psi_x | P | \Psi_x \rangle = \| P | \Psi_x \rangle \|^2$$

where $|\Psi_x\rangle = U(x)|q_0\rangle = U_{x_n} U_{x_{n-1}} ... U_{x_1} |q_0\rangle$.

2.2 Definition of MM-QFA

A measure-many quantum finite automaton is defined by a 6-tuple

$$M = (Q, \Sigma, \delta, q_0, Q_{acc}, Q_{rej})$$

where Q is a finite set of states, Σ is a finite input alphabet with an end-marker symbol \$, δ is a unitary transition function of the same form as for an MO-QFA, and the state q_0 is the initial configuration of M. The set Q is partitioned into three subsets:

Q_{acc} is the set of halting accepting states, Q_{rej} is the set of halting rejecting states, and Q_{non} is the set of non-halting states.

The operation of an MM-QFA is similar to that of an MO-QFA except that after every transition M measures its configuration with respect to the three subspaces that correspond to the three subsets Q_{non}, Q_{acc}, and Q_{rej}: $E_{non} = Span(\{|q\rangle \mid q \in Q_{non}\})$, $E_{acc} = Span(\{|q\rangle \mid q \in Q_{acc}\})$, and $E_{rej} = Span(\{|q\rangle \mid q \in Q_{rej}\})$. If the configuration of M is in E_{non} then the computation continues; if the configuration is in E_{acc} then M accepts, otherwise it rejects. After every measurement the superposition collapses into the measured subspace and is renormalized.

Just like MO-QFAs, the configuration of an MM-QFA is represented by a complex n-dimensional vector, the transition function is represented by unitary matrices, and measurement is represented by diagonal zero-one projection matrices that project the vector onto the respective subspaces.

The definition of an MM-QFA is almost identical to the definition by Kondacs and Watrous in[KW97]. The only difference is that we only require one end-marker at the end of the tape, rather than two end-markers, at the start and end of the tape; this does not affect the acceptance power of the automaton; see Appendix A for further details.

Since M can have a non-zero probability of halting part-way through the computation, it is useful to keep track of the cumulative accepting and rejecting probabilities. Therefore, in some cases we use the representation, of Kondacs and Watrous[KW97] that represents the state of M as a triple $(|\Psi\rangle, p_{acc}, p_{rej})$, where p_{acc} and p_{rej} are the cumulative probabilities of accepting and rejecting. The evolution of M on reading symbol σ is denoted by

$$(P_{non}|\Psi'\rangle,\ p_{acc} + \|P_{acc}|\Psi'\rangle\|^2,\ p_{rej} + \|P_{rej}|\Psi'\rangle\|^2)$$

where $|\Psi'\rangle = U_\sigma|\Psi\rangle$, and P_{acc}, P_{rej}, and P_{non} are the diagonal zero-one projection matrices that project the configuration onto the non-halting, accepting and rejecting subspaces.

2.3 Language Acceptance

A QFA M is said to accept a language L with cut-point λ if for all $x \in L$ the probability of M accepting x is greater than λ and for all $x \notin L$ the probability of M accepting x is at most λ. A QFA M accepts L with bounded error if there exists an $\epsilon > 0$ such that for all $x \in L$ the probability of M accepting x is greater than $\lambda + \epsilon$ and for all $x \notin L$ the probability of M accepting x is less than $\lambda - \epsilon$. We call ϵ the margin.

We partition the languages accepted by QFAs into several natural classes. Let the class $\mathbf{RMO}_\epsilon$ be the set of languages accepted by an MO-QFA with margin of at least ϵ. Let the restricted class of languages, $\mathbf{RMO} = \cup_{\epsilon>0}\mathbf{RMO}_\epsilon$, be the set of languages accepted by an MO-QFA with bounded error, and let the unrestricted class of languages, $\mathbf{UMO} = \mathbf{RMO}_0$, be the set of languages accepted by an MO-QFA with unbounded error. We define the languages classes $\mathbf{RMM}_\epsilon$, **RMM** and **UMM** accepted by an MM-QFA in a similar fashion.

Since the cut-point of a QFA can be arbitrarily raised or lowered, we could without loss of generality fix the cut-point to be $\frac{1}{2}$. However, for the purposes of presentation we use the general cut-point definition stated above.

2.4 Reversible Finite Automata

Unitary operations are reversible, thus QFAs bear strong resemblance to various variants of reversible finite automata. A group finite automaton (GFA) is a deterministic finite automata (DFA) $M = (Q, \Sigma, \delta, q_0, F)$ with the restriction that for every state $q \in Q$ and every input symbol $\sigma \in \Sigma$ there exists exactly one state $q' \in Q$ such that $\delta(q', \sigma) = q$, i.e. δ is a complete one-to-one function and the automaton derived from M by reversing all transitions is deterministic.

A reversible finite automata (RFA) is a DFA $M = (Q, \Sigma, \delta, q_0, F)$ such that for every state $q \in Q$ and for every symbol $\sigma \in \Sigma$ there is at most one state $q' \in Q$ such that $\delta(q', \sigma) = q$, or, if there exist distinct states $q_1, q_2 \in Q$ and symbol $\sigma \in \Sigma$ such that $\delta(q_1, \sigma) = q = \delta(q_2, \sigma)$, then $\delta(q, \Sigma) = \{q\}$. This definition is equivalent to the one used by Ambainis and Freivalds[AF98] and is an extension of Pin's[Pin87] definition.

2.5 Previous Work

Moore and Crutchfield[MC00] introduced a variant of the MO-QFA model and investigated the model in terms of quantum regular languages (QRLs). They showed several closure properties including closure under inverse homomorphisms and derived a method for bilinearizing the representation of an MO-QFA that transforms it into a generalized stochastic system.

Kondacs and Watrous[KW97] introduced a variant of the MM-QFA that was derived from their 2QFA model. Using a technique similar to Rabin's[Rab63], Kondacs and Watrous proved that 1-way QFAs that accept with bounded error are restricted to accepting a proper subset of the regular languages and that the language $L = \{a, b\}^* b$ is not a member of that subset.

Ambainis and Freivalds[AF98] showed that MM-QFAs could accept languages with probability higher than $\frac{7}{9}$ if and only if the language could be accepted by an RFA, which is equivalent to being accepted with certainty by an MM-QFA. In [ABFK99] Ambainis, Bonner, Freivalds, and Kikusts, construct a hierarchy of languages such that the ith language in the hierarchy can be accepted by a MM-QFA with at most probability p_i, where the series (p_i) converges to $\frac{1}{2}$ and is strictly decreasing.

Ambainis, Nayak, Ta-Shma, and Vazirani[ANTSV99], and Nayak[Nay99], investigated how efficiently MM-QFAs can be constructed compared to DFAs. They showed that for some languages the accepting MM-QFA is exponentially larger than the corresponding DFA.

In [AI99] Amano and Iwama studied a restricted version of the 2QFA model where the head was not allowed to move right. They showed that the emptiness problem for this model is undecidable. This is another instance where quantum mechanics provides computational power that is not achievable in the classical case.

3 MO-QFAs

3.1 Bounded Error Acceptance

The restriction that MO-QFAs accept with bounded error is as limiting as in the case of PFAs[Rab63]. Since MM-QFAs can only accept a proper subset of the regular

languages if they are required to accept with bounded error and since every MO-QFA can be simulated exactly by an MM-QFA, the class **RMO** is a proper subset of the regular languages. The class **RMO** is exactly the class of languages accepted by group finite automata (GFAs), otherwise known as group languages, and whose syntactic semigroups are groups, see Eilenberg[Eil76]. This result is implied by Theorem 7 in [MC00] but is not stated in the paper. To prove this result we first need Lemma 3.1.

Lemma 3.1 *Let U be a unitary matrix. For any $\epsilon > 0$ there exists an integer $n > 0$ such that for all vectors x, where $\|x\|^2 \leq 1$, it is true that $\|(I - U^n)x\|^2 < \epsilon$.*

Proof: Let $m = dim(U)$. Since U is a normal matrix, U^n can be written as

$$U^n = PD^nP^{-1}$$

where P is a unitary matrix and D is the diagonal matrix of eigenvalues with the jth eigenvalue having the form $e^{i\pi r_j}$[Ort87]. If all eigenvalues in D are rotations through rational fractions of π, i.e. r_j is rational, then let $n = 2\prod_{j=1}^{m} q_j$ where q_j is the denominator of r_j. Thus $D^n = I$ and we are done.

Otherwise, at least one eigenvalue is a rotation of unity through an irrational fraction of π. Let $l \leq m$ be the number of these eigenvalues. For the other $m-l$ eigenvalues compute n, just as above, and let $D' = D^n$. The value of the jth element on the diagonal of D' is either 1 or $e^{i\pi n r_j}$ where r_j is some irrational real number. Consider taking D' to some power $k \in \mathbb{Z}^+$. The values that are 1 do not change, but the other l values that are of the form $e^{i\theta_j k}$ where $\theta_j = \pi n r_j$, form a vector that varies through a dense subset in an l-dimensional torus. Hence, there exists k such that the l-dimensional vector is arbitrarily close to $\vec{1}$. Thus, for any $\epsilon' > 0$ there exists a $k > 0$ such that $\|(I - D'^k)\vec{1}\|^2 < \epsilon'$. Hence

$$\begin{aligned}
\|(I - U^{nk})x\|^2 &= \|(I - PD^{nk}P^{-1})x\|^2 \\
&= \|P(I - D'^k)P^{-1}x\|^2 \\
&\leq \|(I - D'^k)m\vec{1}\|^2 \\
&= m^2\|(I - D'^k)\vec{1}\|^2 \\
&\leq m^2\epsilon'.
\end{aligned}$$

Select ϵ' such that $\epsilon' < \frac{\epsilon}{m^2}$ to complete the proof. ■

Lemma 3.2, due to Bernstein and Vazirani[BV97], states that if two configurations are close, then the differences in probability distributions of the configurations is small. This lemma relates the closeness of configurations to the variation distance between their probability distributions and allows us to partition the set of reachable configurations into equivalence classes. The variation distance between two probability distributions is the maximum difference in the probabilities of the same event occurring with respect to both distributions.

Lemma 3.2 (Bernstein and Vazirani, 1997)
Let $|\psi\rangle$ and $|\varphi\rangle$ be two complex vector such that $\| |\psi\rangle\|^2 = \| |\varphi\rangle\|^2 = 1$ and $\| |\psi\rangle - |\varphi\rangle\|^2 < \epsilon$. The total variation distance between the probability distributions resulting from measurement of $|\psi\rangle$ and $|\varphi\rangle$ is at most 4ϵ.

Theorem 3.3 follows from these two lemmas.

Theorem 3.3 *A language L can be accepted by an MO-QFA with bounded error if and only if it can be accepted by a GFA.*

Proof: The 'if' direction follows from the fact that the transition function for a GFA is also a valid transition function for an MO-QFA that can accept the same language with certainty.

For the 'only if' direction, by contradiction, assume that there exists a language L that can be accepted by an MO-QFA with bounded error but cannot be accepted by a GFA. Since the class **RMO** is a subset of the regular languages, L must be regular. Let $M = (Q, \Sigma, \delta, q_0, F)$ be an MO-QFA that accepts L with bounded error. If two strings x and y take M into the same reachable configuration, then for any string z the probability of M accepting xz is equal to the probability of M accepting yz, which means that $xz \in L$ if and only if $yz \in L$. Therefore, the space of reachable configurations of M's computation can be partitioned into a finite number of equivalence classes defined by the corresponding minimal DFA for L.

Let $|\psi\rangle$ and $|\varphi\rangle$ denote reachable configurations of M and let $\sim_L$ be the right invariant equivalence relation induced by L. Since L cannot be accepted by a GFA, there must exist two distinct equivalence classes $[y]$ and $[y']$, an equivalence class $[x]$, and a symbol $\sigma \in \Sigma$, such that $[y\sigma] \sim_L [y'\sigma] \sim_L [x]$. If U_σ is the transition matrix for symbol σ, $|\psi\rangle \in [y]$ and $|\varphi\rangle \in [y']$ then $U_\sigma|\psi\rangle \in [x]$ and $U_\sigma|\varphi\rangle \in [x]$.

Since M accepts L with bounded error, let ϵ be the margin. By Lemma 3.1 there exists an integer $k > 0$ such that $\|(I - U_\sigma^k)|\psi\rangle\|^2 < \frac{\epsilon}{4}$ and $\|(I - U_\sigma^k)|\varphi\rangle\|^2 < \frac{\epsilon}{4}$. Hence, $U_\sigma^k|\psi\rangle \in [y]$ because if

$$\begin{aligned} \|(I - U_\sigma^k)|\psi\rangle\|^2 &= \| \, |\psi\rangle - U_\sigma^k|\psi\rangle\|^2 \\ &= \|V(|\psi\rangle - U_\sigma^k|\psi\rangle)\|^2 \\ &< \frac{\epsilon}{4} \end{aligned}$$

where V is an arbitrary unitary matrix, then by Lemma 3.2 the probability of $VU_\sigma^k|\psi\rangle$ being measured in a particular state is within ϵ of $V|\psi\rangle$ being measured in the same state; this probability is less than the margin. Similarly $U_\sigma^k|\varphi\rangle \in [y']$. Hence $[y] \sim_L [y\sigma^k]$ and $[y'] \sim_L [y'\sigma^k]$.

We assumed that $[x] \sim_L [y\sigma] \sim_L [y'\sigma]$ and showed that $[y] \sim_L [y\sigma^k]$ and $[y'] \sim_L [y'\sigma^k]$; therefore, $[y] \sim_L [x\sigma^{k-1}] \sim_L [y']$. Let z be the string that distinguishes $[y]$ and $[y']$. Then the string $\sigma^{k-1}z$ partitions $[x]$ into at least two distinct equivalence classes, but this is a contradiction. Therefore, there cannot exist a language L that can be accepted by an MO-QFA with bounded error but not by a GFA. ■

Theorem 3.3 implies that $\mathbf{RMO}_\epsilon = \mathbf{RMO}_{\epsilon'}$ for all $\epsilon, \epsilon' > 0$, hence there are most two distinct classes of languages accepted by MO-QFAs, the restricted class **RMO**, which is equivalent to the class of languages accepted by a GFA, and the unrestricted class **UMO**.

It follows immediately from Theorem 3.3 that the class **RMO** is closed under boolean operations, inverse homomorphisms, and word quotients, and is not closed under homomorphisms. Moore and Crutchfield[MC00] proved these closure properties for the class **UMO**.

3.2 Non-Regular Languages

Unlike the class **RMO**, the class **UMO** contains languages that are non-regular. This is not surprising given that Rabin[Rab63] proved a similar result for PFAs. In fact our proof closely mimics Rabin's[Rab63] technique.

Lemma 3.4 *Let $L = \{x \in \{a,b\}^* \mid |x|_a \neq |x|_b\}$, there exists a 2-state MO-QFA M that accepts L with cut-point 0.*

Proof: Let $M = (Q, \Sigma, \delta, q_0, F)$ where $Q = \{q_0, q_1\}$, $\Sigma = \{a, b\}$, $F = \{q_1\}$, and δ is defined by the transition matrices

$$U_a = U_b^{-1} = \begin{bmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{bmatrix}$$

where α is an irrational fraction of π. Since U_a is a rotation matrix and α is an irrational fraction of π, the orbit formed by applying U_a to $|q_0\rangle$ is dense in the circle, and there exists only one k, such that $U_a^k|q_0\rangle = |q_0\rangle$, namely $k = 0$. This also holds for $U_b = U_a^{-1}$. Thus, $U(x)|q_0\rangle = |q_0\rangle$ if and only if the number of U_a rotations applied to $|q_0\rangle$ is equal to the number of U_b rotations, which is true if and only if the $|x|_a = |x|_b$. Otherwise, M has a non-zero probability of halting in state q_1. ■

Lemma 3.4 implies that the class **RMO** is properly contained within the class **UMO** and therefore the two classes are distinct.

The MO-QFA in Lemma 3.4 solves the word problem for the infinite cyclic group: is the input word equal to the identity element in the group, where the group has only one generator element, say a, and its inverse $b = a^{-1}$. We can generalize this result to the general word problem for the free group. The word problem for a free group is to decide whether or not a product of a sequence of elements of the free group reduces to the identity[LZ77].

Lemma 3.5 *The word problem for the free group language can be accepted by an MO-QFA.*

Proof: Construct a free group of rotation matrices drawn from the group SO_3 as discussed by Wagon[Wag85]. Let $M = (Q, \Sigma, \delta, q_0, F)$ be a 3-state MO-QFA where $\Sigma = \{a, a^{-1}, b, b^{-1}, ...\}$ such that $|\Sigma|$ is equal to the sum of the number of rotation matrices and their inverses, δ is defined by the rotation matrices and their inverses, and $F = \{q_0\}$. The MO-QFA will accept identity words with certainty and reject non-identity words with a strictly non-zero probability, hence solving the word problem for the free group. ■

3.3 Equivalence of MO-QFAs

In classical automata theory there is an algorithm to determine if two automata are equivalent. We say that QFAs M and M' are equivalent if their probability distributions over Σ^* are the same: for every word $x \in \Sigma$, the probability of M accepting x is equal to the probability of M' accepting x. In order to determine if two MO-QFAs are equivalent we first bilinearize them using the method detailed by Moore and Crutchfield[MC00]; this yields two generalized stochastic systems. We then apply Paz's[Paz71, Page 21, Page 140] method for testing stochastic system equivalence to the generalized stochastic systems to determine if they have the same distribution.

3.4 Simulation of MO-QFAs by PFAs

Most classical computation is either deterministic or probabilistic, hence it is useful to ask how probabilistic automata compare to their quantum analogs. In the case of MO-QFAs, any language accepted by an MO-QFA can also be accepted by a PFA. If L can be accepted by an MO-QFA with bounded error, then it can also be accepted by a PFA with bounded error.

Theorem 3.6 *Let M be an MO-QFA that accepts L with cut-point λ then:*

1. *There exists a PFA that accepts L with some cut-point λ'.*
2. *If M accepts L with bounded error, then there exists a PFA that accepts L with bounded error.*

Proof: The second result follows from Theorem 3.3 because every GFA is also a PFA.

Since we can bilinearize M, L is a generalized cut-point event (GCE)[Paz71, Page 153]. Since the class of GCEs is equal to the class of probabilistic cut-point events (PCEs)[Paz71, Page 153], which are accepted by PFAs, there exists a PFA that can accept L with some cut-point λ'. ■

Combining Theorem 3.6 with Lemma 3.5 yields a new insight into the languages accepted by PFAs:

Corollary 3.7 *The word problem for the free group language can be solved by a PFA.*

4 MM-QFAs

Measure-many quantum finite automata are more powerful than MO-QFAs because a measurement is performed after every transition. This allows the machine to terminate before reading the entire string and simulate the spin states of RFAs.

As mentioned before, an MM-QFA uses one end-marker while the Kondacs and Watrous [KW97] 1-way QFA uses two end-markers. The second marker does not add any more power to the model, see Appendix A, but makes constructing an MM-QFA easier because the MM-QFA can start in an arbitrary configuration. Hence, for the sake of conciseness and clarity we shall assume that some of the MM-QFAs constructed in the following proofs have two end-markers.

4.1 Closure Properties

Unlike the closure properties of the classes **RMO** and **UMO**, which can be derived easily, the closure properties of the classes **RMM** and **UMM** are not as evident and in one important case unknown. We show that the classes **RMM** and **UMM** are closed under complement, inverse homomorphism and word quotient. Similar to the class **RMO**, the class **RMM** is not closed under homomorphisms. It remains an open problem to determine whether the classes **RMM** and **UMM** are closed under boolean operations.

Theorem 4.1 proves that both classes are closed under complement and inverse homomorphisms by showing that each class $\mathbf{RMM}_\epsilon$ is closed under complement and inverse homomorphisms; closure under word quotient follows directly from the latter, given the presence of end-markers.

Theorem 4.1 *The class* $\mathbf{RMM}_\epsilon$ *is closed under complement, inverse homomorphisms, and word quotient.*

Proof: Closure under complement follows from the fact that we can exchange the accept and reject states of the MM-QFA. This exchanges the probabilities of acceptance and rejection but does not affect the margin.

Given an MM-QFA M and a homomorphism h we construct an MM-QFA M' that accepts $h^{-1}(L)$. Let $M = (Q, \Sigma, \delta, Q_{acc}, Q_{rej})$ and $M' = (Q', \Sigma, \delta', Q'_{acc}, Q'_{rej})$. Assume that δ and δ' are defined in terms of matrices $\{U_\sigma\}_{\sigma\in\Sigma}$ and $\{U'_\sigma\}_{\sigma\in\Sigma}$. Unlike the proof for MO-QFAs in [MC00], the direct construction of

$$U'_\sigma = U(h(\sigma))$$

will not work because a measurement occurs between transitions, and combining transitions without taking this into account could produce incorrect configurations. After every transition some amount of probability amplitude is placed in the halting states and should not be allowed to interact with the non-halting states in the following transitions. This is achieved by storing the amplitude in additional states; this technique is also used in [ANTSV99]. Assume without loss of generality that

$$\begin{aligned} Q_{non} &= \{q_i \in Q \mid 0 \le i < n_{non}\} \\ Q_{halt} &= \{q_i \in Q \mid n_{non} \le i < n\} \end{aligned}$$

where $n = |Q|$ and $n_{non} = |Q_{non}|$. Let $m = \max_{\sigma\in\Sigma}\{|h(\sigma)|\}$ and let

$$Q' = Q \cup Q'_{halt}$$

where

$$\begin{aligned} Q'_{halt} &= \{q_i\}_{i=n+1}^{n+m(n-n_{non})} \\ Q'_{acc} &= Q_{acc} \cup \{q_{n+j(i-n_{non})} \in Q'_{halt} \mid q_i \in Q_{acc}, 1 \le j \le m\} \\ Q'_{rej} &= Q_{rej} \cup \{q_{n+j(i-n_{non})} \in Q'_{halt} \mid q_i \in Q_{rej}, 1 \le j \le m\}. \end{aligned}$$

Intuitively, we replicate the halting states m times; each replication is termed a halting state set.

We construct δ' from the matrices of δ. Let V_σ be a unitary block matrix

$$V_\sigma = U_{shift} \begin{bmatrix} U_\sigma & \\ & I_{m(n-n_{non})} \end{bmatrix}$$

where

$$U_{shift} = \begin{bmatrix} I_{n_{non}} & & \\ & & I_{n-n_{non}} \\ & I_{m(n-n_{non})} & \end{bmatrix}.$$

The matrix U_{shift} is a unitary matrix that shifts the amplitudes in the halting set i to the halting set $i+1$ and the amplitude in halting set m to halting set 0. In analogy to the MO-QFA case where $U'_\sigma = U(h(\sigma))$, for MM-QFAs let

$$U'_\sigma = V(h(\sigma)) = V_{x_k} V_{x_{k-1}} ... V_{x_1}$$

where $h(\sigma) = x = x_1x_2...x_k$ and $k \leq m$.

After every x_i sub-transition the halting amplitude is shifted and stored in the $m+1$ halting sets of states. When the sub-transition is done, the amplitude in halt state set 0 is zero, which is what is required to prevent unwanted interactions. A minimum of m sub-transitions must occur before halting set m contains non-zero amplitude, but no more than m sub-transitions will ever occur; therefore halting set 0 will never receive non-zero amplitude from halting set m. Since M' has the same distribution as M, the margin will not decrease.

Closure under word quotient follows from closure under inverse homomorphism and the presence of both end-markers. ■

Just like the class **RMO**, the class **RMM** is not closed under homomorphisms.

Theorem 4.2 *The class* **RMM** *is not closed under homomorphisms.*

Proof: Let $L = \{a, b\}^*c$ and define a homomorphism h to be $h(a) = a$, $h(b) = b$, and $h(c) = b$. Since L can be accepted by an RFA, $L \in$ **RMM** [AF98], but $h(L) = \{a, b\}^*b \notin$ **RMM**, the result follows. ■

A more interesting question is whether the classes **RMM** and **UMM** are closed under boolean operations. Unlike MO-QFAs that have two types of states: accept and reject, MM-QFAs have three types of states: accept, reject, and non-halt. Consequently, the standard procedure of taking the tensor product of two automata to obtain their intersection or union does not work. A general method of intersecting two MM-QFAs is not known. Thus, it is not known whether **RMM** and **UMM** are closed under boolean operations.

4.2 Bounded Error Acceptance

The restriction of bounded error acceptance reduces the class of languages that an MM-QFA can accept to a proper subclass of the regular languages[KW97]. To study the languages in class **RMM**, we look at their corresponding minimal automata. Ambainis and Freivalds[AF98] showed that if the minimal DFA $M(L) = (Q, \Sigma, \delta, q_0, F)$ contains an irreversible construction, defined by two distinct states $q_1, q_2 \in Q$ and strings $x, y, z \in \Sigma^*$ such that $\delta(q_1, x) = \delta(q_2, x) = q_2$, $\delta(q_2, y) \in F$ and $\delta(q_2, z) \notin F$, then an RFA cannot accept L and an MM-QFA cannot accept it with a probability greater than $\frac{7}{9}$; this condition is both sufficient and necessary.

We derive a similar necessary condition for a language L to be a member of the class **RMM**. This condition, called the partial order condition, is a relaxed version of a condition defined by Meyer and Thompson[MT69]. A language L is said to satisfy the partial order condition if the minimal DFA for L satisfies the partial order condition. A DFA satisfies the partial order condition if it does not contain two distinguishable states $q_1, q_2 \in Q$ such that there exists strings $x, y \in \Sigma^+$ where $\delta(q_1, x) = \delta(q_2, x) = q_2$, and $\delta(q_2, y) = q_1$. States q_1 and q_2 are said to be distinguishable if there exists a string $z \in \Sigma^*$ such that $\delta(q_1, z) \in F$ and $\delta(q_2, z) \notin F$ or vice versa[HU79]. Using a result in [KW97], Theorem 4.3 proves that the partial order condition is necessary for an MM-QFA to accept L with bounded error.

Theorem 4.3 *If $M = (Q, \Sigma, \delta, q_0, F)$ is a minimal DFA for language L that does not satisfy the partial order condition then $L \notin$* **RMM**.

Proof: By contradiction, assume that $L \in \mathbf{RMM}$. Let $L_b = \{a,b\}^*b$. Since the minimal DFA for L does not satisfy the partial order condition there exist states $q_1, q_2 \in Q$ and strings $x, y \in \Sigma^+$ as defined above and a distinguishing string $z \in \Sigma^*$ such that $\delta(q_1, z) \notin F$ if and only if $\delta(q_2, z) \in F$. Without loss of generality assume that $\delta(q_1, z) \notin F$ and $\delta(q_2, z) \in F$.

Let s be the shortest string such that $\delta(q_0, s) = q_1$. Let $L' = s^{-1}Lz^{-1}$. By Theorem 4.1, $L' \in \mathbf{RMM}$. Define the homomorphism h as

$$\begin{aligned} h(a) &= xy \\ h(b) &= x \\ h(\Sigma - \{a,b\}) &= xy, \end{aligned}$$

where the last definition is for completeness. Let $L'' = h^{-1}(L')$. By Theorem 4.1 $L'' \in \mathbf{RMM}$. But $L'' = L_b \notin \mathbf{RMM}$, a contradiction. ■

The partial order condition is so named because once the state q_2 is visited, there is no path back to state q_1. Thus, there exists a partial order on the states of the DFA. We do not know whether this condition is also sufficient for MM-QFA acceptance with bounded error. While we do not know whether the class **RMM** is closed under boolean operations, Theorem 4.6 relates closure under intersection to the partial order condition.

Lemma 4.4 *Let M be a DFA that satisfies the partial order condition. The minimal DFA M' that accepts $L(M)$ satisfies the partial order condition.*

Proof: Let $M = (Q, \Sigma, \delta, q_0, F)$ be a DFA and $M' = (Q', \Sigma, \delta', q_0', F')$ be the corresponding minimal DFA. Assume by contradiction that M' does not satisfy the partial order condition. Hence, M' has two states that correspond to the equivalence classes $[q_1']$ and $[q_2']$ such that $[q_1']x \sim_L [q_2']x \sim_L [q_2']$ and $[q_2']y \sim_L [q_1']$. By the Myhill-Nerode theorem[HU79], the equivalence classes partition the set of reachable states in Q. Hence, for each equivalence class $[q_i']$ there is a corresponding subset of Q. Let Q_1 and Q_2 denote the subsets of Q corresponding to the equivalence classes $[q_1']$ and $[q_2']$ and assign an arbitrary order on each subset. Select the first state, say $p_1 \in Q_1$, and define the set $R = \{q \in Q_2 \mid \exists n, m \in \mathbb{Z}^+,\ \delta(p_1, x^m) = \delta(q, x^n) = q\}$. If there exists a state $r \in R$ and string $y \in \Sigma^+$ such that $\delta(r, y) = p_1$, then M does not satisfy the partial order condition, and this is a contradiction. Otherwise, there does not exist a $y \in \Sigma^+$ such that $\delta(r, y) = p_1$ for all $r \in R$. In this case there is a partial order on p_1 and on $Q_1 \backslash \{p_1\}$ because p_1 will never be visited again if M reads a sufficient number of xs. Remove p_1 from Q_1 and repeat the procedure on $p_2 \in Q_1$. After a finite number of iterations we will either find a p_i that satisfies our requirements, which means that M does not satisfy the partial order condition and is a contradiction, or none of the states in Q_1 will have the required characteristics, in which case M' satisfies the partial order condition. Therefore, if M satisfies the partial order condition, so will its minimal equivalent M'. ■

Lemma 4.5 *Let L' and L'' be languages that satisfy the partial order condition. Then $L = L' \cap L''$ also satisfies the partial order condition.*

Proof: Let $M' = (Q', \Sigma, \delta', q_0', F')$ be the minimal DFA accepting the language L' and let $M'' = (Q'', \Sigma, \delta'', q_0'', F'')$ be the minimal DFA accepting the language L''. We first construct an automaton M that accepts $L' \cap L''$ by combining M' and M'' using a direct product. Define $M = (Q, \Sigma, \delta, q_{00}, F)$ where $Q = Q' \times Q''$, $q_{00} = (q_0', q_0'')$, $F = \{(q', q'') \in Q \mid q' \in F' \ \wedge \ q'' \in F''\}$ and $\delta((q', q''), \sigma) = (\delta'(q', \sigma), \delta''(q'', \sigma))$.

We argue that if M' and M'' satisfy the partial order condition, then so will M. Assume, by contradiction, that M does not satisfy the partial order condition. Then there exist two states $q_{ij} = (q_i', q_j'')$ and $q_{kl} = (q_k', q_l'')$ and strings $x, y, z \in \Sigma^+$ such that $\delta(q_{ij}, x) = \delta(q_{kl}, x) = q_{kl}$, $\delta(q_{kl}, y) = q_{ij}$ and $\delta(q_{ij}, z) \in F$ if and only if $\delta(q_{kl}, z) \notin F$. In the first case assume that either $i \neq k$ or $j \neq l$, and without loss of generality, assume the former. Then there exists state $q_i' \in Q'$ and state $q_k' \in Q'$ such that $\delta'(q_i', x) = \delta'(q_k', x) = q_k'$, $\delta_1(q_k', y) = q_i'$. But this means that M' does not satisfy the partial order condition, a contradiction. In the second case assume that $i = k$ and $j = l$. This implies that $q_{ij} = q_{kl}$ and hence there cannot exist a string z that distinguishes the two states, also a contradiction. Therefore M must satisfy the condition.

Since M satisfies the partial order condition and accepts L, by Lemma 4.4 the minimal automaton that accepts L satisfies the partial order condition, and hence L itself, satisfies the partial order condition. ■

Theorem 4.6 *If the partial order condition is sufficient for acceptance with bounded error by MM-QFAs then the class* **RMM** *is closed under intersection.*

Proof: By Lemma 4.5 the intersection of two languages that satisfy the partial order condition is a language that satisfies the partial order condition. ■

One method for proving that the class **RMM** is not closed under intersection involves intersecting two languages in **RMM** and showing that the resulting language is not in **RMM**. By Theorem 4.6 this method will not work unless the partial order condition is insufficient. To study whether the partial order condition is sufficient, as well as necessary, we show that a well known class of languages can be accepted by an MM-QFA with bounded error.

4.3 Piecewise Testable Sets

A piecewise testable set is a boolean combination of sets of the form

$$L_x = \Sigma^* x_1 \Sigma^* x_2 \Sigma^* ... \Sigma^* x_n \Sigma_*$$

where $x_i \in \Sigma$ [Per94]. Intuitively, L_x is the language of strings that contain the successive symbols of x as a subsequence; we call such a language a partial piecewise testable set.

Piecewise testable sets, introduced by Simon in [Sim75], form a natural family of star-free languages. Such sets define a class of computations that wait for a partially ordered sequence of trigger events (input symbols); if a trigger event (symbol) is read that is not next in the sequence, it is simply ignored. Another natural interpretation of piecewise testable sets is subsequence searching. Consider a language where a word is said to be in the language if it contains a finite boolean combination of subsequences. Such a language is a piecewise testable language and word acceptance corresponds to

searching the words for the required subsequences. Finally, such languages belong to a class of languages whose MM-QFAs have an arbitrarily large, but finite, set of ordered states.

We show, by construction, that MM-QFAs can accept partial piecewise testable sets with bounded error. Furthermore, the MM-QFAs we construct are of a special kind that we call 'end-decisive' and the languages accepted by end-decisive MM-QFAs with bounded error are closed under intersection, which implies that MM-QFAs can accept piecewise testable sets with bounded error. To construct these MM-QFAs we introduce several useful concepts: junk states, end-decisiveness, and trigger chains.

A junk state is a halting state of a MM-QFA in which the probability of the machine accepting is fixed to cut-point λ and the probability of rejecting is fixed at $1-\lambda$; it is implemented by two halting states, accept and reject. Any probability amplitude flowing into a junk state is split between an accepting and rejecting state. That is, if α amplitude flows into a junk state, then $\sqrt{\lambda}\alpha$ amplitude flows into an accept state, and $\sqrt{1-\lambda}\alpha$ flows into a reject state. While junk states are implemented using standard accept and reject states, we treat the junk state as a separate halting state. Any accept or reject state that is not used to implement a junk state is called a decisive state. Consequently, we say that a specifically designed MM-QFAs may halt in a decisive state, accepting or rejecting, or in a junk state. Intuitively, a junk state signals a failed computation.

A machine is end-decisive if it halts in a decisive state only after reading the end-marker of the input string. The machine may halt before reading the end-marker but may do so only in a junk state. Intuitively, a computation that ends before the entire string is read is treated as a failed computation. Formally, a machine is end-decisive if the only transition that allows probability amplitude to flow into a decisive halting state is the \$ transition.

An end-decisive machine that accepts with bounded error has probability, bounded by some constant $\tau < 1$ of ending up in a junk state and a probability $1-\tau$ of ending up in a decisive state. If $\tau \not< 1$ then the amount of probability amplitude ending up in a decisive small can become arbitrary small, dropping below any fixed margin. Thus, τ must be strictly less than one for the MM-QFA to accept with bounded error.

The probability of an end-decisive machine accepting is the sum of the probability of the machine ending up in a non-decisive accepting state and the probability of the machine ending up in a decisive accepting state. The probability of machine M accepting string x is therefore

$$\begin{aligned} P[M(x) = \text{accept}] &= \tau P[M(x) = \text{non}-\text{decisive accept} \mid \text{failure}] + \\ &\quad (1-\tau)P[M(x) = \text{decisive accept} \mid \text{success}] \\ &= \tau\lambda + (1-\tau)P[M(x) = \text{accept} \mid \text{success}]. \end{aligned}$$

In order for M to accept L with cut-point λ and margin ϵ, it must hold that

$$P[M(x) = \text{accept} \mid \text{success}] > \lambda + \mu$$

if $x \in L$ and

$$P[M(x) = \text{accept} \mid \text{success}] < \lambda - \mu$$

if $x \notin L$, where $\mu = \frac{\epsilon}{1-\tau}$ is called the decisive margin. Therefore, the cut-point is based on the decisive states of the machine. The margin, in part, is dictated by the decisive states also.

A trigger chain is a construction of junk states and transition matrices that causes a reduction in amplitude of a particular state only if the amplitude of another state is decreased, presumably by some previous transition. Trigger chains correspond directly to partial piecewise testable sets. Consider the matrix

$$X = \begin{bmatrix} \frac{1}{2} & \frac{1}{\sqrt{2}} & \frac{1}{2} \\ \frac{1}{\sqrt{2}} & 0 & -\frac{1}{\sqrt{2}} \\ \frac{1}{2} & -\frac{1}{\sqrt{2}} & \frac{1}{2} \end{bmatrix}.$$

This matrix is a special case of a transition matrix introduced by Ambainis and Freivalds[AF98]. This matrix operates on three states and is a triggering mechanism of the chain. Consider the vectors

$$|\psi\rangle = (\alpha, 0, \beta)^T$$

and

$$X|\psi\rangle = \left(\frac{\alpha}{2} + \frac{\beta}{2}, \frac{\alpha}{\sqrt{2}} - \frac{\beta}{\sqrt{2}}, \frac{\alpha}{2} + \frac{\beta}{2}\right)^T.$$

The vectors $|\psi\rangle$ and $X|\psi\rangle$ are equal if and only if $\alpha = \beta$. If $\alpha \neq \beta$ then the amplitudes of the first and third state are averaged, with the remainder of the amplitude going into the second state. We define a generalized version of X by embedding it into a larger identity block matrix. Define V_i to be

$$X_i = \begin{bmatrix} I_i & & \\ & X & \\ & & I_{n-i-3} \end{bmatrix}$$

where I_m is an $m \times m$ identity matrix, X is defined as above, and n is the number of states, i.e. the size of X_i. The matrix X_i operates on a triple of states, q_i through to q_{i+2}. We assume that state q_{i+1}, the second state, is a junk state unless otherwise noted. Each junk state is implemented by an accepting and a rejecting state, but for the purposes of clarity we do not discuss them.

Theorem 4.7 *Let L_x be a partial piecewise testable language. There exists an end-decisive MM-QFA that accepts L_x with bounded error.*

Proof: We construct an MM-QFA M that accepts L_x where $x = x_0x_1...x_n$.

For each link in the trigger chain we require a junk state and a non-halting state. We order the states to correspond with the description of the X_i matrices. Specifically, the first $2n + 2$ states are the non-halting states, interleaved with junk states. Each triple of states $(q_2i, q_{2i+1}, q_{2i+2})$ corresponds to a link of the trigger chain, of which there are $n+1$. State q_{2n+1} is the decisive accept state and the remaining $2n+2$ states are used to implement the junk states.

Let $m = 4n + 8$ and $M = (Q, \Sigma, \delta, q_0, Q_{acc}, Q_{rej})$ where

$$\begin{aligned} Q &= \{q_0, ..., q_m\} \\ Q_{junk} &= \{q_i \in Q \mid 0 < i < 2n \ \wedge \ i \equiv 1 \bmod 2\} \cup \{q_{2n+3}, q_{2n+4}\} \\ Q_{acc} &= \{q_i \in Q \mid 2n + 4 < i \leq 3n + 6\} \cup \{q_{2n+1}\} \\ Q_{rej} &= \{q_i \in Q \mid 3n + 6 < i \leq 4n + 8\} \end{aligned}$$

and the last $2n + 4$ states are used to implement the junk states.

Define δ by the transition matrices $\{U_\sigma\}_{\sigma \in \Sigma}$. Each transition matrix U_σ consists of a product of matrices:

$$U_\sigma = JU_{\sigma,n}U_{\sigma,n-1}...U_{\sigma,0}$$

where

$$U_{\sigma,i} = \begin{cases} S & i = 0 \ \wedge \ x_0 = \sigma \\ X_{2i-2} & 1 \leq i \leq n \ \wedge \ x_i = \sigma \\ I_m & \text{otherwise} \end{cases}$$

and

$$S = \begin{bmatrix} 0 & 1 & \\ 1 & 0 & \\ & & I_{m-2} \end{bmatrix}.$$

The matrix S shifts the amplitude of q_0 to the junk state q_1. This is the first trigger that is activated when x_0 is read. The matrix J implements the junk states; it transfers the amplitude from the junk states to the accepting and rejecting states. Finally, let the transition matrix for the end-marker \$ be

$$U_\$ = JFX_{2n}$$

where

$$F = \begin{bmatrix} R & & & & & & & & \\ & \ddots & & & & & & & \\ & & R & & & & & & \\ & & & 0 & 0 & 0 & 0 & 1 & \\ & & & 0 & 1 & 0 & 0 & 0 & \\ & & & 0 & 0 & 0 & 1 & 0 & \\ & & & 0 & 0 & 1 & 0 & 0 & \\ & & & 1 & 0 & 0 & 0 & 0 & \\ & & & & & & & & I_{2n+4} \end{bmatrix}$$

and the matrix

$$R = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

The matrix F sends all amplitude into the junk states. The matrix X_{2n} sends some minimum amount of amplitude into an accept state if the amplitudes of states q_{2n} and q_{2n+2} differ.

The initial configuration of the machine is $|\psi_{init}\rangle = (\alpha_0, \alpha_1, ..., \alpha_{2n+4})^T$ where

$$\alpha_i = \begin{cases} \frac{1}{\sqrt{n+2}} & 0 \leq i \leq 2n+2 \ \wedge \ i \equiv 0 \pmod 2 \\ 0 & \text{otherwise} \end{cases}$$

i.e. the amplitude is evenly distributed among all non-halting states.

The only decisive accepting state in the machine is q_{2n+1}, and amplitude only flows into it when the end-marker is read. In order for it to get a non-zero amplitude, the amplitudes of states q_{2n} and q_{2n+2} must differ. Since all non-halting states start with the same amplitude, and since the amplitude of state q_{2n+2} will not change during the execution of the machine until the end-marker is read, the amplitude of state q_{2n-2} must change in order for the amplitude of state q_{2n} to change. Following the same argument, state q_{2i} will not change in amplitude, until state q_{2i-2} changes in amplitude. Furthermore, the change in amplitude of state q_{2i} is governed by the matrix components X_{2i-2} and X_{2i}. Hence, the initial change of amplitude of state q_{2i} depends exclusively on a change in amplitude of state q_{2i-2} and is governed by component X_{2i-2} that is located in the transition matrix U_{x_i}. If any other transition matrix is applied, then the amplitude of state q_{2i} will not change. Hence M can read $(\Sigma - \{x_i\})^*$ without changing the amplitude of state q_{2i}, but, as soon as x_i is read, component X_{2i-2} will be applied and q_{2i} will have a decreased amplitude, provided state q_{2i-2} already had a decrease of its amplitude. Finally, the amplitude of any state q_{2i} will never increase beyond its initial value, and once the amplitude of state q_{2i} decreases, it will never increase beyond $\frac{1}{\sqrt{n+2}} - (\frac{1}{2})^{n+2}$. For the case of symbol x_0, the amplitude of state q_0 is changed by matrix S to 0 and is the starting trigger. When the end-marker is read a minimum of $\frac{1}{\sqrt{n+2}}(\frac{1}{2})^{n+3}$ of amplitude is placed into the accepting state only if the amplitude of state q_{2n} has decreased. The rest of the amplitude, from all $n+2$ non-halting states is channeled into junk states. If the amplitudes of q_{2n} and q_{2n+2} do not differ then all amplitude is channeled into junk states.

The probability of M accepting a string not in the language is $\frac{1}{2}$, while the probability of M accepting a string in the language is at least $\frac{1}{2} + \frac{1}{n+2}(\frac{1}{2})^{2n+7}$. If we want to bound the probability of rejecting strings that are not in the language away from $\frac{1}{2}$, we add two more states to M: a non-halting state that holds a small amount of amplitude, and a reject state into which the amplitude gets dumped upon reading the \$ symbol. The amount of amplitude required is a some small fraction, ϵ', of $\sqrt{\frac{1}{n+2}(\frac{1}{2})^{2n+7}}$. Consequently, the margin is $\epsilon'^2 \frac{1}{n+2}(\frac{1}{2})^{2n+7}$. Therefore, M accepts with bounded error. ■

The MM-QFA construction in Theorem 4.7 is end-decisive. This means that we can take intersections and complements of partial piecewise testable languages.

Lemma 4.8 *Let M be an end-decisive MM-QFA that accepts with bounded error. The complement of M is also an end-decisive MM-QFA that accepts with bounded error.*

Proof: Construct a new MM-QFA from M by leaving the junk states as they are and exchanging all the decisive accepting and rejecting states of M. The new automaton is end-decisive and has the same margin as M. ■

Lemma 4.9 *Let M and M' be end-decisive MM-QFAs that accept L and L' respectively, with bounded error. There exists an end-decisive MM-QFA M'' that accepts $L'' = L \cap L'$ with bounded error.*

Proof: To facilitate the proof we use a slightly different formalism to describe MM-QFAs. Let $M = (Q, \Sigma, \delta, q_0, Q_{acc}, Q_{rej}, Q_{junk})$ and $M' = (Q', \Sigma, \delta', q'_0, Q'_{acc}, Q'_{rej}, Q'_{junk})$ be end-decisive MM-QFAs that accept L and L'. Using these two MM-QFAs we construct an MM-QFA $M'' = (Q'', \Sigma, \delta'', q''_0, Q''_{acc}, Q''_{rej}, Q''_{junk})$ that accepts $L'' = L \cap L'$. In this case Q_{acc} and Q_{rej} are sets that contain decisive accept and reject states, while Q_{junk} contains junk halting states, each of which can be implemented by an accepting and rejecting state. This formalism is equivalent to the standard one.

First, ignore the non-decisive accept and reject states of the two MM-QFAs. That is treat Q and Q' as being composed of only non-halting states, decisive halting states, and junk states. As the last step in our construction we implement the junk states in M''

Second, we construct M''. Let $Q'' = Q \times Q'$ and $q''_0 = (q_0, q'_0)$. The three sets of halting states are defined as

$$\begin{aligned} Q''_{junk} &= \{(q_i, q'_j) \in Q'' \mid q_i \in Q_{junk} \ \vee \ q'_j \in Q'_{junk}\} \\ Q''_{acc} &= \{(q_i, q'_j) \in Q'' \mid q_i \in Q_{acc} \ \wedge \ q'_j \in Q'_{acc}\} \\ Q''_{rej} &= \{(q_i, q'_j) \in Q'' \mid (q_i, q'_j) \notin Q''_{junk} \ \wedge \ (q_i \in Q_{rej} \ \vee \ q'_j \in Q'_{rej})\} \end{aligned}$$

and the transition function δ'' is defined as

$$\delta''((q, q'), \sigma, (r, r')) = \delta(q, \sigma, r) \cdot \delta'(q', \sigma, r'),$$

which is a tensor product of the transition functions δ and δ', ignoring the non-decisive accept and reject states.

Third, we implement the junk states of M'' by adding an additional non-decisive accept and reject state for each junk state $q'' \in Q''_{junk}$. We also modify the transitions such that the probability amplitude from the junk states flows into the accept and reject states with amplitude $\sqrt{\lambda''}$ and amplitude $\sqrt{1 - \lambda''}$, where λ'' will be the cut-point of the new MM-QFA.

Since M and M' are end-decisive, i.e., the decisive states will only have non-zero amplitude when the end-marker is read, the MM-QFA M'' will be end-decisive. We only need to derive a cut-point λ'' and margin ϵ'' for M''.

If $x \in L''$, then $P[M(x) = \text{accept} \mid \text{success}] > \lambda + \mu$ and $P[M'(x) = \text{accept} \mid \text{success}] = \lambda' + \mu'$, where λ and λ' are the cut-points with which M and M' accept L and L' and, μ and μ' are the decisive margins. By summing over the decisive accepting states it follows that the probability of M'' accepting string x given that M'' completes the computation on string x, i.e., enters a deciding state, is the product of the accepting probabilities of $M(x)$ and $M'(x)$ given that M and M' enter decisive states;

$$\begin{aligned} P[M''(x) = \text{accept} \mid \text{success}] &= P[M(x) = \text{accept} \mid \text{success}] \cdot \\ &\quad P[M'(x) = \text{accept} \mid \text{success}]. \end{aligned}$$

If $x \in L''$ then

$$\begin{aligned} P[M''(x) = \text{accept} \mid \text{success}] &= P[M(x) = \text{accept} \mid \text{success}] \cdot \\ &\quad P[M'(x) = \text{accept} \mid \text{success}] \\ &> (\lambda + \mu) \cdot (\lambda' + \mu') \\ &= \lambda\lambda' + \lambda'\mu + \lambda\mu' + \mu\mu' \\ &= (\lambda\lambda' + \lambda'\mu) + (\lambda\mu' + \mu\mu'). \end{aligned}$$

On the other hand, assume without loss of generality that $\lambda\mu' \leq \lambda'\mu$; if $x \notin L''$ then

$$\begin{aligned} P[M''(x) = \text{accept} \mid \text{success}] &= P[M(x) = \text{accept} \mid \text{success}] \cdot \\ &\quad P[M'(x) = \text{accept} \mid \text{success}] \\ &< (\lambda - \mu) \cdot (\lambda' + \mu') \\ &= \lambda\lambda' + \lambda'\mu - \lambda\mu' - \mu\mu' \\ &= (\lambda\lambda' + \lambda'\mu) - (\lambda\mu' + \mu\mu'). \end{aligned}$$

Let $\lambda'' = \lambda\lambda' + \lambda'\mu$, $\mu'' = \lambda\mu' + \mu\mu'$, and implement the junk states in M'' with the cut-point λ''. Since both M and M' accept L and L' respectively, with bounded error, the probability of computation failure of $M''(x)$ is bounded by some constant τ. Consequently, the margin of M'' becomes $\epsilon'' = \mu'' \cdot (1 - \tau) > 0$. Therefore, L'' can also be accepted with bounded error. Thus, M'' accepts L'' with bounded error, which implies that class of languages accepted by end-decisive MM-QFAs with bounded error is closed under intersection. ■

Corollary 4.10 *Let M and M' be end-decisive MM-QFAs that accept with bounded error. There exists an end-decisive MM-QFA M'' that accepts $L = L(M) \cup L(M')$ with bounded error.*

Theorem 4.11 *Piecewise testable sets can be accepted by MM-QFAs with bounded error.*

5 Conclusions

We defined two models of 1-way quantum finite automata: the measure-once model that performs one measurement at the end of the computation, and the measure-many model that performs a measurement after every transition. The measure-many model is strictly more powerful than the measure-once but is more difficult to characterize.

When restricted to accepting with bounded error, measure-once automata can only accept group languages, while unrestricted measure-once automata can accept irregular sets and in particular, can solve the word problem on the free group. Any language accepted by a MO-QFA can also be accepted by a PFA, therefore PFAs can also solve

the word problem on the free group. We also sketched an algorithm for determining equivalence of two MO-QFAs.

The measure-many automaton is difficult to characterize. We have shown that the two classes of languages, those accepted with and without bounded error, are closed under complement and inverse homomorphisms; it is still an open question if these classes are closed under boolean operations. We defined the partial order condition for languages and proved that it is a necessary condition for a language to be accepted by an MM-QFA with bounded error. We also showed that piecewise testable sets can be accepted with bounded error by MM-QFAs, and in the process detailed several novel construction techniques.

We do not know if the partial order condition is also a sufficient condition for bounded acceptance. If it is then the two classes of languages accepted by an MM-QFA are closed under intersection.

References

[ABFK99] A. Ambainis, R. Bonner, R. Freivalds, and A. Kikusts. Probabilities to accept languages by quantum finite automata. In *Computation and Combinatorics*, volume 1627 of *Lecture Notes on Computer Science*, 1999.

[AF98] A. Ambainis and R. Freivalds. 1-way quantum finite automata: Strengths, weaknesses and generalizations. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, pages 332–342, November 1998.

[AI99] M. Amano and K. Iwama. Undecideability of quantum finite automata. In *Proceedings of the 31st Annual ACM Symposium on the Theory of Computing*, pages 368–375, 1999.

[ANTSV99] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani. Dense quantum coding and a lower bound for 1-way quantum automata. In *Proceedings of the 31st Annual ACM Symposium on the Theory of Computing*, pages 376–383, 1999.

[BV97] E. Bernstein and U. Vazirani. Quantum complexity theory. *SIAM Journal of Computing*, pages 1411–1473, October 1997.

[Eil76] S. Eilenberg. *Automata, Languages and Machines*, volume B. Academic Press, New York, 1976.

[HU79] J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Languages and Computation.* Addison-Wesley Publishers, Reading, Massachusetts, 1979.

[KW97] A. Kondacs and J. Watrous. On the power of quantum finite state automata. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 66–75, 1997.

[LZ77] R. Lipton and Y. Zalcstein. Word problem solvable in logspace. *Journal of the ACM*, 24(3):523–526, July 1977.

[MC00] C. Moore and J. Crutchfield. Quantum automata and quantum grammars. *Theoretical Computer Science*, 237:275–306, 2000.

[MT69] A. Meyer and C. Thompson. Remarks on algebraic decomposition of automata. *Mathematical Systems Theory*, 3(2):110–118, 1969.

[Nay99] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, 1999.

[Ort87] J. Ortega. *Matrix Theory*. Plenum Press, New York, New York, 1987.

[Paz71] A. Paz. *Introduction to Probabilistic Automata*. Academic Press, New York, New York, 1971.

[Per94] D. Perrin. Finite automata. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 1. Elsevier Science Publisher, 1994.

[Pin87] J. Pin. On languages accepted by finite reversible automata. In *Proceedings of the 14th International Colloquium on Automata, Languages and Programming*, volume 267 of *Lecture Notes on Computer Science*, pages 237–249, 1987.

[Rab63] M. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.

[Sim75] I. Simon. Peicewise testable events. In *Proc. of the 2nd GI Conf*, volume 33 of *Lecture Notes on Computer Science*, 1975.

[Wag85] S. Wagon. *The Banach-Tarski Paradox*. Cambridge University Press, New York, New York, 1985.

A End-Marker Theorems

Theorem A.1 *Let M be an MO-QFA that has both left and right end-markers. There exists an MO-QFA M' that uses only one end-marker and is equivalent to M.*

Proof: Let $M = (Q, \Sigma, \delta, q_0, F)$ be an MO-QFA with left and right end-markers, effectively allowing M to start in any possible configuration. Define $M' = (Q, \Sigma, \delta', q_0, F)$ from M. Let δ be defined in terms of the transition matrices $\{U_\sigma\}_{\sigma\in\Sigma}$. We define δ' from δ in the following way: for every $\sigma \in \Sigma$ let

$$U'_\sigma = U_{\not\in}^{-1} U_\sigma U_{\not\in}$$

and let

$$U'_\$ = U_\$ U_{\not\in}$$

Now consider what happens when M and M' read a string $x = x_1...x_n$. Since

$$\begin{aligned} U'(x\$) &= U'_\$ U'_{x_n} ... U'_{x_1} \\ &= U_\$ U_{\not\in} U_{\not\in}^{-1} U_{x_n} U_{\not\in} ... U_{\not\in}^{-1} U_{x_1} U_{\not\in} \\ &= U_\$ U_{x_n} ... U_{x_1} U_{\not\in} \\ &= U(\not\in x\$), \end{aligned}$$

the probability of M accepting x is equal to the probability of M' accepting x. Thus one end-marker on the right suffices, and by symmetry one left end-marker would also suffice. Therefore, an MO-QFA starting in configuration $|q_0\rangle$ can simulate an MO-QFA starting in any arbitrary configuration. ■

Theorem A.2 *Let M be an MM-QFA that has both left and right end-markers. There exists an MM-QFA M' that uses only a right end-marker and is equivalent to M.*

Proof: Let $M = (Q, \Sigma, \delta, Q_{acc}, Q_{rej})$ be an MM-QFA that uses two end-markers and accepts L. Assume without loss of generality that

$$\begin{aligned} Q_{non} &= \{q_i \in Q \mid 0 \leq i < n_{non}\} \\ Q_{acc} &= \{q_i \in Q \mid n_{non} \leq i < n_{acc}\} \\ Q_{rej} &= \{q_i \in Q \mid n_{acc} \leq i < n_{rej} = n = |Q|\}, \end{aligned}$$

which facilitates a simpler description of M'. We construct $M' = (Q', \Sigma, \delta', Q'_{acc}, Q'_{rej})$ that accepts L with only the right end-marker. Let $Q' = Q \cup \{q_n, q_{n+1}, ..., q_{2n-n_{non}}\}$, $Q'_{acc} = \{q_{n+i-n_{non}} \in Q' \mid q_i \in Q_{acc}\}$ and $Q'_{rej} = \{q_{n+i-n_{non}} \in Q' \mid q_i \in Q_{rej}\}$. Assume that δ is defined in terms of transition matrices $\{U_\sigma\}_{\sigma\in\Sigma}$. The construction of $\{U'_\sigma\}_{\sigma\in\Sigma}$ is similar to that in the proof of Theorem A.1. Let I_l represent an identity matrix of size l and $m = n - n_{non}$. We define δ' in terms of its unitary block matrices. For all $\sigma \in \Sigma$ let

$$\begin{aligned} U'_\sigma &= \begin{bmatrix} U_{\not\$}^{-1} & \\ & I_m \end{bmatrix} S \begin{bmatrix} U_\sigma & \\ & I_m \end{bmatrix} \begin{bmatrix} U_{\not\$} & \\ & I_m \end{bmatrix} \\ U'_\$ &= S \begin{bmatrix} U_\$ & \\ & I_m \end{bmatrix} \begin{bmatrix} U_{\not\$} & \\ & I_m \end{bmatrix} \end{aligned}$$

where

$$S = \begin{bmatrix} I_{n_{non}} & & & \\ & & & I_m \\ & & I_m & \end{bmatrix}$$

transfers (sweeps) all probability amplitude from states in the old halting states to the new halting states. The old halting states, those in Q_{acc} and Q_{rej}, are no longer halting states in M'. The operation of M' is similar to the operation of the QFA constructed in Theorem A.1, The "sweeping" operation saves the amplitude that was in the old states, while it performs the $U_{\not\$}^{-1}$ operation in the new halting states (since otherwise the $U_{\not\$}^{-1}$ would corrupt the amplitude stored in the original halting states). ■