Electronic Colloquium on Computational Complexity, Report No. 159 (2014)

Magic coins are useful for small-space quantum machines

A. C. Cem Say
Boğaziçi University, Department of Computer Engineering, Bebek 34342 İstanbul, Turkey
Abuzer Yakaryılmaz*
National Laboratory for Scientific Computing, Petrópolis, RJ, 25651-075, Brazil
say@boun.edu.tr abuzer@lncc.br

Keywords: quantum computation, constant space, unrestricted amplitudes

Abstract

Although polynomial-time probabilistic Turing machines can utilize uncomputable transition probabilities to recognize uncountably many languages with bounded error when allowed to use logarithmic space, it is known that such "magic coins" give no additional computational power to constant-space versions of those machines. We show that adding a few quantum bits to the model changes the picture dramatically. For every language L, there exists such a two-way quantum finite automaton that recognizes a language of the same Turing degree as L with bounded error in polynomial time. When used as verifiers in public-coin interactive proof systems, such automata can verify membership in all languages with bounded error, outperforming their classical counterparts, which are known to fail for the palindromes language.

1 Introduction

Determining whether various resource-bounded quantum models are equivalent or superior to their classical probabilistic counterparts in power gives valuable insight about the nature of quantum computation, and about the important cases (notably, polynomial-time models) where this problem is still open. Sometimes, arguably unrealistic setups and resources (e.g. [1, 3]) are also used in these comparisons, keeping in mind how the manifestly unrealistic concept of nondeterminism forms a priceless tool of complexity theory. In this paper, we will be considering variants of bounded-error constant-space Turing machines allowed to toss coins with unrestricted (i.e. possibly uncomputable) real numbers as the probability of coming up heads. In that regime, it is known that probabilistic polynomial-time language recognizers have no additional computational power over their standard, fair-coin versions [8]. Languages that cannot be handled by probabilistic public-coin verifiers, even with such magic coins, have also been demonstrated [7]. We show that quantum computers outperform their classical counterparts in these contexts;

*Yakaryılmaz was partially supported by CAPES with grant 88881.030338/2013-01, ERC Advanced Grant MQC, and FP7 FET project QALGO. Moreover, the part of the research work was done while Yakaryılmaz was visiting Boğaziçi University in 2014.

ISSN 1433-8092

by demonstrating that every language is Turing-equivalent to some language recognized by a polynomial-time two-way quantum finite automaton, and that every language has a public-coin proof system with such an automaton as the verifier. The amount of "quantumness" used in our constructions is very small, as the machines we utilize can be viewed as deterministic two-way finite automata augmented by just one or two qubits.

2 Preliminaries

It is well known that polynomial-time probabilistic Turing machines (PTM's) can utilize uncomputable transition probabilities to recognize uncountably many languages with bounded error: The k'th bit in the decimal expansion of the probability that a given coin will land heads can be estimated by a procedure that involves tossing that coin for a number of times that is exponential in k. Given any language L on the alphabet $\{a\}$, and a coin which lands heads with probability $0.x$, where x is an infinite sequence of digits whose k'th member encodes whether the k'th unary string is in L, the language $\{a^{4^k} | a^k \in L\}$ is therefore in $\mathsf{BPP}_{\mathbb{R}}$, the version of BPP with arbitrary real transition probabilities.[1] The machine in this construction uses logarithmic space, and it was proven by Dwork and Stockmeyer [8] that no polynomial-time PTM using $o(\log \log n)$ space can recognize a nonregular language with bounded error, even with unrestricted transition probabilities.

Adleman et al. [4] used a construction similar to the one described above to demonstrate that $\mathsf{BQP}_{\mathbb{R}}$, the class of languages recognized with bounded error in polynomial time by quantum Turing machines (QTM's) employing arbitrary real amplitudes, contains languages of every Turing degree. The QTM formulated by Adleman et al. also uses logarithmic space. Our construction in Section 3 shows that the coins described in [4] can also be exploited by polynomial-time constant-space PTM's augmented by a single qubit.

The constant-space quantum model that we will use is the two-way finite automaton with quantum and classical states (2qcfa), introduced by Ambainis and Watrous [6], in which the quantum and classical memories are nicely separated, allowing a precise quantification of the amount of "quantumness" required for the task at hand.[2] These machines can be viewed as two-way deterministic automata augmented with a quantum register of constant size (typically, just one or two qubits). The input string is assumed to be written between two end-marker symbols on a tape. A 2qcfa starts with both its quantum and classical parts set to be in their respective initial states, and operates separately on these parts at each step:

- First, a superoperator (Figure 1), determined by the current classical state and the symbol being scanned on the input tape, is applied to the quantum register, yielding an outcome.
- Then, the next classical state and tape head movement direction is determined by the current classical state, the symbol being scanned on the input tape, and the observed outcome.

[1] We thank Peter Shor for helping us with these facts.

[2] Our exposition of 2qcfa's will use superoperators, generalizing and simplifying the setup of [6]; see [17]. It is not hard to modify our algorithms for implementation in the format of [6].

Execution halts when the machine enters an accepting or rejecting classical state.

For a 2qcfa with j quantum states, each superoperator $\mathcal{E}$ is composed of a finite number of $j \times j$ matrices called *operation elements*, $\mathcal{E} = \{E_1, \ldots, E_k\}$, satisfying

$$\sum_{i=1}^{k} E_i^\dagger E_i = I, \tag{1}$$

where $k \in \mathbb{Z}^+$, and the indices label the possible outcomes. When a superoperator $\mathcal{E}$ is applied to a quantum register in superposition $|\psi\rangle$, one observes the outcome i with probability $p_i = \langle\widetilde{\psi_i}|\widetilde{\psi_i}\rangle$, where $|\widetilde{\psi_i}\rangle$ is calculated as $|\widetilde{\psi_i}\rangle = E_i|\psi\rangle$, and $1 \leq i \leq k$. If the outcome i is observed ($p_i > 0$), the new superposition is obtained by normalizing $|\widetilde{\psi_i}\rangle$, yielding $|\psi_i\rangle = \frac{|\widetilde{\psi_i}\rangle}{\sqrt{p_i}}$. Unitary operations such as rotations, as well as measurements and probabilistic branchings can be realized within this framework.

Figure 1: Superoperators (adapted from [18])

Ambainis and Watrous [6] demonstrated a 2qcfa with just two quantum states (i.e. a single qubit) and computable transition amplitudes that recognizes the nonregular language $\texttt{EQ} = \{a^n b^n \mid n \geq 0\}$, in polynomial expected time, thereby establishing the superiority of such machines over their classical counterparts. In the next section, we will use a modified version of the Ambainis-Watrous algorithm as part of our construction of 2qcfa's of the same Turing degree as any given language.

We will compare the capabilities of probabilistic and quantum finite automata as verifiers in public-coin proof systems in Section 4. It is known that $\mathsf{AM}_{\mathbb{R}}(\mathsf{2pfa})$, the class of languages which have such bounded-error proof systems with two-way probabilistic finite automata utilizing arbitrary real transition probabilities as verifiers, does not include the binary palindromes language [7].

In public-coin systems with 2qcfa verifiers [14], the outcomes of the superoperators of the verifier are instantly available to the computationally powerful prover, who is trying to convince the verifier that the input should be accepted. The quantum and classical transition functions of the 2qcfa's in these systems are modified to take into account of the communication symbol sent by the prover in each step; see [18] for the details.

Although this paper is the first to study 2qcfa verifiers with unrestricted amplitudes, the class $\mathsf{AM}_{\mathbb{Q}}(\mathsf{2qcfa})$, i.e. the version with rational amplitudes, has already been shown to include PSPACE[3] and some NEXP-complete languages [14]. The constructions in this papers also involve very small amounts of quantum memory. One technique that we will borrow from [14, 18] is that a 2qcfa verifier can encode the integer represented by a binary string that is read from the prover into the amplitude of a quantum state, albeit with low probability in each such attempt.

Another demonstration of the superiority of quantum finite automata over their classical counterparts in this unrealistic setup with no noise and decoherence is the result of Aaronson and Drucker [2], who showed a succinctness advantage of quantum real-time machines in the task of distinguishing two classical coins with different bias.

In the following, $\Sigma^*(i)$ denotes the ith element in the lexicographic ordering of all possible strings on alphabet Σ, where $i > 0$. $\Sigma^*(1)$ is the empty string, denoted ε. We

[3]The proof will appear in the new version of [14].

will use the fact that, for any i, the string $1\Sigma^*(i)$ is the binary encoding of i. For a string w, $|w|_\sigma$ denotes the number of occurrences of symbol σ in w.

3 Language recognition in polynomial time

In this section, we present a constant-space version of a theorem proven by Adleman et al. [4] for logarithmic-space quantum Turing machines. The main concern in the adaptation is to get a machine to recognize, and count up to, arbitrary powers of 8, without using any secondary memory.

Theorem 1. *For every language L, there exists a language L' on the alphabet $\{a, b\}$ such that L' is Turing equivalent to L, and there exists a 2qcfa which recognizes L' with bounded error in polynomial expected time.*

Proof. We start by considering the language

$$\texttt{POWER-EQ} = \{aba^7ba^{7\cdot 8}ba^{7\cdot 8^2}ba^{7\cdot 8^3}b\cdots ba^{7\cdot 8^n} \mid n \geq 0\},$$

where the number of a's in any string is seen to be 8^{n+1} for some $n \geq 0$. The 2qcfa described below (Figure 2) uses a single qubit (with quantum state set $\{q_0, q_1\}$) to recognize POWER-EQ in polynomial time.

If the input is aba^7, ACCEPT.
Check whether the input is of the form $aba^7ba^{7\cdot t_1}ba^{7\cdot t_2}\cdots ba^{7\cdot t_n}$ for some $n > 0$, where all the t_i are positive multiples of 8. If not, REJECT.
Move the head to the leftmost b in the input.
LOOP:
  Move the head right to the next a.
  Set the qubit to $|q_0\rangle$.
  While the currently scanned symbol is a:
    Rotate the qubit with angle $\sqrt{2}\pi$.
    Move the head to the right.
  (The currently scanned symbol is a b.) Move the head to the a on the right.
  While the currently scanned symbol is a:
    Rotate the qubit with angle $-\sqrt{2}\pi$.
    Move the head 8 squares to the right.
  Measure the qubit. If the result is q_1, REJECT. (I)
  If the currently scanned symbol is a b, move the head to the nearest b on the left, and goto LOOP.
  (The currently scanned symbol is the right end-marker.)
  Repeat twice:
    Move the tape head to the first input symbol.
    While the currently scanned symbol is not an end-marker, do the following:
      Simulate a classical coin flip. If the result is heads, move right. Otherwise, move left.
  If the process ended at the right end-marker both times, and two more coin flips both turn out heads, goto EXIT.
  Move the head to the leftmost b in the input, and goto LOOP.
EXIT:
  ACCEPT the input.

Figure 2: A 2qcfa for POWER-EQ

After an easy deterministic check, that 2qcfa enters an infinite loop where each iteration compares t_i with $\frac{t_{i+1}}{8}$ for all $i \in 1, \ldots, n-1$. This is achieved by first rotating the qubit counterclockwise t_i times, then rotating it clockwise $\frac{t_{i+1}}{8}$ times, and finally checking whether it has returned to its original orientation $|q_0\rangle$. Since the rotation angle is an irrational multiple of π, the probability r_i that the machine will reject at the line marked (I) in Figure 2 is zero if and only if $t_i = \frac{t_{i+1}}{8}$ for the corresponding i. If $t_i \neq \frac{t_{i+1}}{8}$, then r_i will be at least [6]

$$\frac{1}{2(t_i - \frac{t_{i+1}}{8})^2} > \frac{1}{2(t_i + \frac{t_{i+1}}{8})^2}.$$

We therefore conclude that any input string $w \notin$ POWER-EQ which has survived the deterministic check in the beginning will be rejected with a probability greater than $\frac{1}{2|w|^2}$ in each iteration of the infinite loop.

If the input w has not been rejected after all the $n-1$ comparisons described above, the 2qcfa makes two consecutive random walks starting on the first input symbol, and ending at an end-marker. The probability that both these walks will end at the right end-marker, leading to acceptance in this iteration of the infinite loop, is $\frac{1}{(|w|+1)^2}$, and the expected runtime for this stage is $O(|w|^2)$ [6]. This means that the machine will halt within $O(|w|^2)$ expected iterations of the loop, leading to an overall expected runtime of $O(|w|^4)$.

To conclude, the 2qcfa of Figure 2 will accept any string $w \in$ POWER-EQ with probability 1. On the other hand, any $w \notin$ POWER-EQ that makes it into the loop has a rejection probability that is more than twice as large as its acceptance probability in each iteration, and therefore will be rejected with probability greater than $\frac{2}{3}$.

We are now ready to adapt the technique of Adleman et al. [4] to 2qcfa's, since the algorithm in Figure 2 provides us with a way of ensuring that the number of a's on the tape is a power of 8 with sufficiently high reliability.

For any $\mathtt{L} \subseteq \Sigma^*$, define the language

$$\texttt{POWER-EQ}(\mathtt{L}) = \{w \in \{a,b\}^* \mid w \in \texttt{POWER-EQ} \text{ and } \Sigma^*(\log_8(|w|_a)) \in \mathtt{L}\}.$$

Note that L and POWER-EQ(L) are Turing reducible to each other.

We will show that, for any L, POWER-EQ(L) can be recognized by the 2qcfa described in Figure 3 in polynomial expected time:

(Assume that the input string is in POWER-EQ.)
Move the head to the left end-marker.
Set the qubit to $|q_0\rangle$.
While the currently scanned symbol is not the right end-marker:
  Rotate the qubit with angle $\theta_{\mathtt{L}}$ only if the currently scanned symbol is an a.
  Move the head to the next square on the right.
Rotate the qubit with angle $\frac{\pi}{4}$.
Measure the qubit. If the result is q_1, ACCEPT. Otherwise, REJECT.

Figure 3: A 2qcfa for POWER-EQ(L)

Any string which is a member of POWER-EQ contains 8^j a's for some j, and that string

is in POWER-EQ(L) if $\Sigma^*(j) \in \mathtt{L}$. This information about L is encoded into the transition amplitudes via the angle $\theta_{\mathtt{L}}$ as follows.

Define

$$\theta_{\mathtt{L}} = 2\pi \sum_{i=1}^{\infty} \left(\frac{F_{\mathtt{L}}(i)}{8^{i+1}} \right)$$

where the function $F_{\mathtt{L}} : \mathbb{N} \to \{-1, 1\}$ is

$$F_{\mathtt{L}}(n) = \begin{cases} 1, & \text{if } \Sigma^*(n) \in \mathtt{L}, \\ -1, & \text{if } \Sigma^*(n) \notin \mathtt{L}. \end{cases}$$

Rotating the qubit 8^j times with $\theta_{\mathtt{L}}$ radians leaves it at an angle

$$8^j \cdot 2\pi \sum_{i=1}^{\infty} \left(\frac{F_{\mathtt{L}}(i)}{8^{i+1}} \right) = \pi \frac{F_{\mathtt{L}}(j)}{4} + \left(2\pi \frac{F_{\mathtt{L}}(j+1)}{8^2} + 2\pi \frac{F_{\mathtt{L}}(j+2)}{8^3} + \cdots \right) \pmod{2\pi}$$

radians from the original orientation $|q_0\rangle$. After the last rotation by $\frac{\pi}{4}$ radians, the qubit's final angle from $|q_0\rangle$ is $\frac{\pi}{2} + \delta$ (i.e. near $|q_1\rangle$) if $F_{\mathtt{L}}(j) = 1$, and δ (i.e. near $|q_0\rangle$) if $F_{\mathtt{L}}(j) = -1$ for a δ guaranteed to be sufficiently small to obtain an error bound of 0.02, as in [4]. The runtime is linear in the input size.

All that remains is to combine the algorithms of Figures 2 and 3. One starts by using the technique depicted in Figure 2 to reject strings that are not in POWER-EQ with probability at least 0.666. Members of POWER-EQ are treated correctly by the 2qcfa of Figure 3 with probability 0.98. We conclude that the combined 2qcfa recognizes POWER-EQ(L) with probability at least 0.65, and this probability can be reduced further as desired using standard repetition techniques.

The expected running time of the combined algorithm is mainly due to the method of Figure 2, and is again polynomially bounded. □

Some automata-theoretic corollaries of this result are presented in the Appendix.

4 Public-coin proof systems

In our demonstration of the power provided to verifiers by quantum coins with uncomputable bias, we will use a slightly different method than the one in Section 3 to represent an entire language by a real number.

For any given language L on alphabet Σ, let the characteristic function $G_{\mathtt{L}}$ map any member of L to 1, and any non-member to 0. The number

$$\gamma_{\mathtt{L}} = \sum_{i=1}^{\infty} \frac{G_{\mathtt{L}}(\Sigma^*(i))}{4^i} = \frac{G_{\mathtt{L}}(\Sigma^*(1))}{4} + \frac{G_{\mathtt{L}}(\Sigma^*(2))}{4^2} + \frac{G_{\mathtt{L}}(\Sigma^*(3))}{4^3} + \cdots$$

will be used to encode information about L into the bias of a quantum coin, as will be explained shortly. Note that $\gamma_{\mathtt{L}}$ equals 0 for $\mathtt{L} = \{\}$, and $\frac{1}{3}$ for $\mathtt{L} = \Sigma^*$. For any other language L, $\gamma_{\mathtt{L}}$ is a real number between 0 and $\frac{1}{3}$.

We start with the case of tally languages, which allows a simpler illustration of the main technique.

Theorem 2. *For any language* L *on the unary alphabet* $U = \{a\}$*, there exists a bounded-error public-coin interactive proof system where the messages are classical, the verifier is a 2qcfa with a single quantum bit, and the expected runtime for inputs of length* n *is* $2^{O(n)}$*.*

Proof. We describe the proof system in question. For any input of length $n \geq 0$, the prover is supposed to send a stream of bits describing the membership status of every string of length at most n, i.e. the sequence

$$u_{\mathtt{L},n} = G_{\mathtt{L}}(\varepsilon)G_{\mathtt{L}}(a)G_{\mathtt{L}}(aa)G_{\mathtt{L}}(aaa)\cdots G_{\mathtt{L}}(a^k)$$

to the verifier. This transmission is controlled by the outcomes of the coins of the verifier, which are of course visible to the prover. As will be seen below, some outcomes will be interpreted as commands for interrupting the transmission and restarting it from the first bit, whereas others will simply mean "go on with the next bit".

Figure 4 depicts the superoperators corresponding to the "coins" of the verifier. The outcome associated with each operation element appears as the subscript in its name. The reader may check that all four superoperators obey the wellformedness condition (1) stated in Figure 1. The initial state of the quantum part is the one corresponding to the first row and column in the operation elements, so the initial superposition of the qubit is just $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$.

The operation elements of $INIT$ are

$$INIT_{go1} = \frac{1}{\sqrt{1+\gamma_{\mathtt{L}}^2}} \begin{pmatrix} \gamma_{\mathtt{L}} & 0 \\ 1 & 0 \end{pmatrix} \text{ and } INIT_{go2} = \frac{1}{\sqrt{1+\gamma_{\mathtt{L}}^2}} \begin{pmatrix} 0 & \gamma_{\mathtt{L}} \\ 0 & 1 \end{pmatrix}.$$

The operation elements of $PROC$-0 are

$$PROC\text{-}0_{go} = \frac{1}{4} \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix} \text{ and } PROC\text{-}0_{restart} = \frac{\sqrt{15}}{4} \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

The operation elements of $PROC$-1 are

$$PROC\text{-}1_{go} = \frac{1}{\sqrt{18}} \begin{pmatrix} 4 & -1 \\ 0 & 1 \end{pmatrix} \text{ and } PROC\text{-}1_{restart} = \frac{1}{\sqrt{18}} \begin{pmatrix} 1 & 4 \\ 1 & 0 \end{pmatrix}.$$

The operation elements of $RIGHT$ are

$$RIGHT_{reject} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, RIGHT_{agree} = \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{\sqrt{3}} \end{pmatrix}, \text{ and } RIGHT_{restart} = \begin{pmatrix} 0 & \frac{\sqrt{2}}{\sqrt{3}} \\ 0 & 0 \end{pmatrix}.$$

Figure 4: Nontrivial superoperators used by the verifier in the proof of Theorem 2

The first action taken by the verifier is an application of the superoperator $INIT$ to the qubit, setting it to the superposition

$$\frac{1}{\sqrt{1+\gamma_{\mathtt{L}}^2}} \begin{pmatrix} \gamma_{\mathtt{L}} \\ 1 \end{pmatrix}. \tag{2}$$

The outcome "*go*" associated with this application cues the prover to send the first bit of $u_{\mathtt{L},n}$, i.e. the membership indicator of the empty string in L. The input head of the verifier is on the left end-marker at this point.

For each bit $b \in \{0,1\}$ sent by the prover, the verifier applies the corresponding superoperator $PROC\text{-}b$ to the qubit. If the outcome turns out to be a "*go*", the verifier moves its input head one square to the right, and the same process is repeated unless the right end-marker is scanned. Any "*restart*" outcome observed at any point causes the verifier to move the head back to the left end-marker and restart the program.

Assume that the first j symbols of the input have been processed in this manner without any restarts occurring. Write the qubit's superposition at that point in the form

$$\alpha \begin{pmatrix} \delta \\ 1 \end{pmatrix}, \tag{3}$$

for two real numbers α and δ, where $\alpha = \frac{1}{\sqrt{1+\delta^2}}$. For $j = 0$, $\delta = \gamma_{\mathtt{L}}$, as stated above. If the prover now sends a 0 (respectively, a 1), and no restart occurs due to the application of the corresponding superoperator, the next superposition would be

$$\beta \begin{pmatrix} 4\delta \\ 1 \end{pmatrix} \text{(respectively, } \eta \begin{pmatrix} 4\delta - 1 \\ 1 \end{pmatrix}, \text{)}$$

for some reals β and η.

Recalling that $\gamma_{\mathtt{L}} \in [0, 1/3]$, it is important to note that, as long as the prover sends membership bits consistent with the encoding of $\mathtt{L}$ in $\gamma_{\mathtt{L}}$, the value of δ in Expression (3) will always be in the interval $[0, 1/3]$ for any $j \leq n$. Otherwise, this value jumps out of the interval $(-2/3, 1)$ at the point of the first disagreement between the prover transmission and $\gamma_{\mathtt{L}}$, and never returns to this interval later in the execution for greater values of j, no matter what the prover says.

If the head reaches the right end-marker, the verifier applies the superoperator $RIGHT$, described in Figure 4 when scanning the right end-marker. The three sides of this "coin" lead to rejection, or a decision to "agree" with the prover (i.e. give the same decision as the prover's claim about the string a^k communicated in the last transmitted bit), or a restart, as indicated in the figure.

As discussed in [16], the overall acceptance probability of such a "program with restart" equals the ratio of the probability of acceptance in a single "round" without any restarts to the total probability of halting in such a round. This simplifies the analysis of the behavior of our verifier.

Prior to the application of the superoperator $RIGHT$, the qubit will again be in a superposition of the form in Expression (3), say,

$$\alpha_{final} \begin{pmatrix} \delta_{final} \\ 1 \end{pmatrix}.$$

The automatic rejection probability due to the operation element $RIGHT_{reject}$ is then $\alpha^2_{final}\delta^2_{final}$, whereas the probability of adopting the prover's decision due to $RIGHT_{agree}$ is $\alpha^2_{final}/3$.

If the input string is a member of $\mathtt{L}$, then in any particular round, a prover obeying the protocol will cause the verifier to accept with probability $\alpha^2_{final}/3$, whereas the rejection probability will be at most $\alpha^2_{final}/9$, since $\delta \leq 1/3$ in that case, as explained above. The overall acceptance probability is therefore at least $3/4$.

If the input a^n is not in $\mathtt{L}$, then the verifier has to transmit an incorrect value for at least the bit about a^n to avoid a rejection probability of 1. But in this case, the absolute value of the amplitude of the first quantum state will be at least $2\alpha_{final}/3$ as explained above, leading to a rejection probability of at least $4\alpha_{final}^2/9$ in a single round. The overall probability of rejection would then be at least 4/7.

To provide a bound on the expected runtime of this algorithm, we examine the probability that the procedure will halt in a single round. Note that both operation elements of the $INIT$ operator leave the qubit in superposition in Expression [2], so the analysis for the first round is the same as the ones which are caused by restarts. All the remaining operators will find the qubit in the form of Expression (3). Scrutiny of Figure 4 shows that an application of $PROC$-0 will yield a "go" with probability at least 1/16, and $PROC$-1 will yield a "go" with more than 0.05 probability at each application. The $RIGHT$ operator ends up with a decision to halt with at least 1/3 probability.[4] This leads to an upper bound of $O(20^n)$ for the expected number of rounds. Since each run takes linear time, we conclude that the verifier runs in time $2^{O(n)}$. □

Theorem 3. *For any language* $\mathtt{L}$ *on the binary alphabet* $B = \{0, 1\}$*, there exists a bounded-error public-coin interactive proof system where the messages are classical, the verifier is a 2qcfa with two quantum bits, and the expected runtime for inputs of length* n *is* $2^{2^{O(n)}}$*.*

Proof. We construct the proof system in question. For an input string w, the transmission expected from a truthful prover will be

$$b_{\mathtt{L},w} = \#\varepsilon\ddagger G_{\mathtt{L}}(\varepsilon)\#0\ddagger G_{\mathtt{L}}(0)\#1\ddagger G_{\mathtt{L}}(1)\#00\ddagger G_{\mathtt{L}}(00)\#01\ddagger G_{\mathtt{L}}(01)\#10\ddagger G_{\mathtt{L}}(10)\cdots\#w\ddagger G_{\mathtt{L}}(w),$$

containing the strings and their membership bits in lexicographic order up to w. (This transmission can be interrupted and restarted by verifier coin outcomes, as we saw in the previous proof.) Note that the first two symbols of b_L are ‡‡, since ε has length 0.

The verifier has four quantum states, to which we will refer with their places in the state vector. It sets its quantum register to the superposition

$$\frac{1}{\sqrt{\gamma_{\mathtt{L}}^2+3}}\begin{pmatrix}\gamma_{\mathtt{L}}\\1\\1\\1\end{pmatrix}$$

at the beginning of the computation, and after every restart, as explained in the proof of Theorem 2. The prover then begins the transmission of what is supposed to be $b_{\mathtt{L},w}$, bit by bit. Figure 5 describes the action of the verifier on each block of the form $\#s\ddagger\sigma$ of the prover message, where $s \in B^*$, and $\sigma \in B$. Note that the procedure depicted in Figure 5 can be interrupted by a "restart" outcome, as in the proof of Theorem 2.

This procedure is implemented by the application of the superoperators described in Figure 6. The operation elements corresponding to restarts, which have been omitted from the figure in the interest of brevity, and the common coefficient c, can be constructed

[4] These bounds are calculated by considering the value of δ that maximizes the restarting probability in each case.

- Perform the following two tasks in parallel while reading the prover's description of the string s:
  - Scan the input string w from left to right, comparing it to s. If s is longer than w, reject.
  - Encode the binary integer $1s$ into the amplitude of the third quantum state.
- When the ‡ symbol, indicating the end of s, is received, compare the amplitude third quantum state with that of the fourth one, rejecting with some probability only if they are different.
- Treat the membership bit σ in the same manner as in Figure 4, computing a new value for the amplitude of the first state. If $s = w$, the decision on the input will be given after the processing of σ at the end of this block, in the same way as in Theorem 2, by executing the superoperator $DECIDE$ described in Figure 6.

Figure 5: Verifier procedure on a segment $\#s \ddagger \sigma$ of the prover message

easily to satisfy Equation 1. It can be seen that, if one starts in a superposition of the form

$$\alpha \begin{pmatrix} \beta \\ 1 \\ 1 \\ K \end{pmatrix},$$

then the application of a sequence of operation elements $ENCODE\text{-}b_{go}$ corresponding to the ordering of bits b in s leave the register in a superposition of the form

$$\iota \begin{pmatrix} \beta \\ 1 \\ N \\ K \end{pmatrix}, \tag{4}$$

where α, β, ζ and K are some real numbers, and N is the integer represented by the binary string $1s$.)

The superoperator $SUCC$ (Figure 6), applied when the input symbol ‡ is scanned, has the effect of rejecting the input with some probability that is nonzero only if $N \neq K$, restarting with some of the remaining probability, and going on with $N + 1$ encoded into the fourth amplitude for being used later in the processing of the next segment otherwise.

The superoperators $PROC$-0 and $PROC$-1 for treating the membership bit are completely analogous to their namesakes in Figure 4.

The $DECIDE$ superoperator is applied only at the end of the final transmission block of each round, where the prover sends the actual input string and its purported membership bit.

It can be seen that the basic idea is the same as in Theorem 2, with the added complication that an evil prover can now trick the finite-state verifier about which string it is reporting a membership bit for. The solution is to have the prover spell out the strings, and the verifier to compare the encoding of each string with that of the previously sent one to try to catch any such tricks.

Let us analyze this algorithm by tracing the unconditional probabilities of acceptance and rejection in a single round: If the input string is in L, the $SUCC$ operator will never

$$ENCODE\text{-}0_{go} = c\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad ENCODE\text{-}1_{go} = c\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$SUCC_{go} = c\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \end{pmatrix} \quad SUCC_{reject} = c\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

$$PROC\text{-}0_{go} = c\begin{pmatrix} 4 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad PROC\text{-}1_{go} = c\begin{pmatrix} 4 & -1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$DECIDE_{reject} = c\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad DECIDE_{agree} = c\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \frac{1}{\sqrt{3}} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

Figure 6: Nontrivial operation elements used by the verifier in the proof of Theorem 3

cause a rejection. The probability that the input is rejected is at most $(c^j \frac{1}{3})^2$, and the acceptance probability is $(c^j \frac{1}{\sqrt{3}})^2$, for an integer j dependent on the input, leading to overall acceptance with probability $\frac{3}{4}$, as in the proof of Theorem 2.

What if the input w is not in $\mathtt{L}$? If the prover sends the strings up to w in correct lexicographic order and attempts to sneak in an incorrect membership bit, the verifier rejects with an overall probability of $\frac{4}{7}$, as analyzed in the proof of Theorem 2. If a string is indeed presented out of order, the $SUCC$ operator will catch this and reject with probability $p = (c^j(N-K))^2$ (cf. Expression (4)) for some $j > 0$. Any acceptance decision that may occur later as a result of the $DECIDE_{agree}$ operation element would necessarily have probability $(c^k \frac{1}{\sqrt{3}})^2$ for some $k > j$, and it is easy to see that this is small in comparison with p, leading to the conclusion that w will be rejected with high overall probability.

Since $b_{\mathtt{L},w}$ has length $2^{O(n)}$, the halting probability in each round is $c^{-2^{O(n)}}$, yielding an expected running time double exponential in n. □

For any decidable language $\mathtt{L}$, $\gamma_{\mathtt{L}}$ is a computable number. Proof systems like the ones described in Theorems 2 and 3, but with only computable amplitudes, therefore exist for all decidable languages.

5 Open questions

The capabilities of bounded error PTM's utilizing $o(\log n)$ space and possibly uncomputable transition probabilities, both as recognizers and verifiers, are unknown. Can they handle uncountably many languages, like their quantum counterparts?

Debate systems [18] are generalizations of proof systems, where an additional agent, the refuter, is trying to convince the verifier that the prover is wrong. Are the classical or quantum versions of these systems able to make better use of uncomputable transitions than the single-prover systems?

Wang [12] showed that the class of all languages recognized by 2pfa's using rational transition probabilities is strictly contained in the class of all deterministic context-sensitive languages. Can 2pfa's with uncomputable transitions recognize any undecidable language with bounded-error? Does the answer change if we allow the machine to use sublogarithmic space (or a counter, or pebbles, or multiple heads, etc.)?

What can be said in this regard about 2pfa's with uncomputable transitions that are verifiers in private-coin proof systems?

Acknowledgement

We thank Peter Shor for his helpful answers to our questions.

References

[1] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A*, 461(2063):3473–3482, 2005.

[2] Scott Aaronson and Andrew Drucker. Advice coins for classical and quantum computation. In *ICALP (1)*, volume 6755 of *LNCS*, pages 61–72. Springer, 2011.

[3] Scott Aaronson and John Watrous. Closed timelike curves make quantum and classical computing equivalent. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2102):631–647, 2009.

[4] Leonard M. Adleman, Jonathan DeMarrais, and Ming-Deh A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, 1997.

[5] Masami Amano and Kazuo Iwama. Undecidability on quantum finite automata. In *STOC'99: Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 368–375, 1999.

[6] Andris Ambainis and John Watrous. Two–way finite automata with quantum and classical states. *Theoretical Computer Science*, 287(1):299–311, 2002.

[7] Cynthia Dwork and Larry Stockmeyer. Finite state verifiers I: The power of interaction. *Journal of the ACM*, 39(4):800–828, 1992.

[8] Cynthia Dwork and Larry J. Stockmeyer. A time complexity gap for two-way probabilistic finite-state automata. *SIAM Journal on Computing*, 19(6):1011–1123, 1990.

[9] Attila Kondacs and John Watrous. On the power of quantum finite state automata. In *FOCS'97*, pages 66–75, 1997.

[10] Michael O. Rabin. Probabilistic automata. *Information and Control*, 6:230–243, 1963.

[11] Arseny M. Shur and Abuzer Yakaryılmaz. Quantum, stochastic, and pseudo stochastic languages with few states. In *UCNC 2014*, volume 8553 of *LNCS*, pages 327–339. Springer, 2014.

[12] Jie Wang. A note on two-way probabilistic automata. *Information Processing Letters*, 43:321–326, 1992.

[13] Abuzer Yakaryılmaz. Superiority of one-way and realtime quantum machines. *RAIRO - Theoretical Informatics and Applications*, 46(4):615–641, 2012.

[14] Abuzer Yakaryılmaz. Public qubits versus private coins. In *The Proceedings of Workshop on Quantum and Classical Complexity*, pages 45–60. Univeristy of Latvia Press, 2013. ECCC:TR12-130.

[15] Abuzer Yakaryılmaz and A. C. Cem Say. Efficient probability amplification in two-way quantum finite automata. *Theoretical Computer Science*, 410(20):1932–1941, 2009.

[16] Abuzer Yakaryılmaz and A. C. Cem Say. Succinctness of two-way probabilistic and quantum finite automata. *Discrete Mathematics and Theoretical Computer Science*, 12(2):19–40, 2010.

[17] Abuzer Yakaryılmaz and A. C. Cem Say. Unbounded-error quantum computation with small space bounds. *Information and Computation*, 279(6):873–892, 2011.

[18] Abuzer Yakaryılmaz, A. C. Cem Say, and H. Gökalp Demirci. Debates with small transparent quantum verifiers. In *Proceedings of the 18th International Conference on Developments in Language Theory*, pages 327–338, 2014.

[19] Shenggen Zheng, Daowen Qiu, Lvzhou Li, and Jozef Gruska. One-way finite automata with quantum and classical states. In *Languages Alive*, volume 7300 of *LNCS*, pages 273–290, 2012.

A Some corollaries of Theorem 1

A language is said to be *stochastic* if there exists a one-way probabilistic finite automaton (1pfa) that accepts any member with probability greater than $\frac{1}{2}$, and rejects any non-member with probability at most $\frac{1}{2}$. In his seminal paper, Rabin proved that the cardinality of the class of stochastic languages is uncountable [10].[5] The constructions in Section 3 allow us to give an alternative proof of this result.

Theorem 4. *There exist uncountably many stochastic languages.*

[5] Rabin's proof uses a binary alphabet. Recently, it was shown that [11] the cardinality of stochastic languages on a unary alphabet is also uncountable.

Proof. We will prove that POWER-EQ(L) is a stochastic language for any given language L. This will complete the proof, since there are uncountably many L, and a different POWER-EQ(L) for any L.

It is known that the class of languages recognized by the one-way versions of 2qcfa's (namely, the 1qcfa's [19]) or any other one-way qfa variant with cut-point $\frac{1}{2}$, i.e. as described in the definition of stochastic languages above, is again the class of stochastic languages [17]. It will therefore be sufficient to show how to construct a 1qcfa algorithm, say M, that accepts all and only the members of POWER-EQ(L) are accepted with probability bigger than $\frac{1}{2}$, for any given L.

Being a one-way machine, M reads the input from left to the right symbol by symbol in a single pass.

While reading the input, M executes some classical and quantum procedures in parallel. There are two classical procedures:

- If the input is aba^7, accept.
- Check the input to see whether it is of the form

$$aba^7ba^{7\cdot t_1}ba^{7\cdot t_2}\cdots ba^{7\cdot t_n} \tag{5}$$

  for some $n > 0$, where all the t_i are positive multiples of 8. If not, reject.

We now describe the quantum procedures, for which M uses four qubits:

1. Use the first qubit to compare t_1 with $\frac{t_2}{8}$, t_3 with $\frac{t_4}{8}$, and so on. To compare t_{2j-1} with $\frac{t_{2j}}{8}$, we set the qubit to $|q_0\rangle$, and then rotate it with angle $\sqrt{2}\pi$ t_{2j-1} times, and with angle $-\sqrt{2}\pi$ $\frac{t_{2j}}{8}$ times. Then measure the qubit, rejecting if q_1 is observed. Otherwise, continue with the next pair.
2. Use the second qubit to apply the same pairwise comparison to t_2 and $\frac{t_3}{8}$, t_4 and $\frac{t_5}{8}$, and so on.
3. Use the third qubit to flip a fair coin for each a that is scanned. If any flip results in a "tail", do not take a decision in this procedure. Use the fourth qubit to implement the procedure given in Figure 3 as long as you keep getting all "head"s on the third qubit: That is, set the fourth qubit to $|q_0\rangle$ at the beginning, and then rotate it with angle θ_L (described in Section 3) for each a you see. At the end, rotate the qubit with angle $\frac{\pi}{4}$ and then make a measurement. If the outcome is q_1, accept. Otherwise, reject.

If any one of these procedures reaches the right end-marker without a decision, it flips a fair coin to accept and reject with equal probability.

If the input is not of the form (5), then it is accepted with probability 0. We assume that the input is of the form (5) for the rest of the discussion. The analysis of the first and second quantum procedures was given in the proof of Theorem 1. Therefore, any input that is not in POWER-EQ is rejected with a polynomially small probability by those procedures. Note that the third quantum procedure will take a decision with exponentially

small probability. A majority of the computational paths of the machine will end up contributing equal amounts to the overall acceptance and rejection probabilities.

Therefore, for any input not in POWER-EQ, the rejection caused by the first and/or second quantum procedures will overwhelm any acceptance due to the third one, with the input being accepted with total probability less than $\frac{1}{2}$. If the input is in POWER-EQ, then whether the overall acceptance probability is greater or less than $\frac{1}{2}$ depends on the tiny but mostly correct contribution from the third procedure. Thus, POWER-EQ(L) is a stochastic language. □

It is still open whether 2qcfa's can recognize any unary non-regular language with bounded error.

We define two more languages: $\texttt{UPOWER} = \{a^{8^n} \mid n \geq 0\}$ and $\texttt{UPOWER}(\texttt{L}) = \{a^{8^n} \mid n \in \texttt{L}\}$. We know that UPOWER can be recognized by a two-way deterministic one-counter automaton (2dca) in $O(|w| \log |w|)$ time, where w is the given input. A 2qcfa augmented with a classical counter is abbreviated as 2qcca.

Theorem 5. *For any given* L, UPOWER(L) *can be recognized by a 2qcca with bounded-error in* $O(|w| \log |w|)$ *time.*

Proof. Let $w = a^m$ be the input, where $m \geq 0$. The 2qcca algorithm consists of two phases. It first classically determines whether $w \in \texttt{UPOWER}$, by checking if m is a power of 8. If so, it executes the quantum phase, which checks for membership in UPOWER(L) by taking the membership function of L into account.

In its first pass over the input w, the classical phase rejects if $m = 0$, or if m is not a multiple of 8, and sets the value of the counter to $\frac{m}{8}$ otherwise.

At this point with the head on the right end-marker, the machine controls the value in the counter, and switches to the quantum phase (described below) if it equals 1. Otherwise, the head is moved exactly $\frac{m}{8}$ symbols from right to left, setting the counter to 0 at the end of this walk. The head then scans the same $\frac{m}{8}$-symbol postfix of the input from left to right, incrementing the counter once every 8 steps, setting its value to $\lfloor \frac{m}{8^2} \rfloor$ when the right end-marker is reached, and checking whether the length of this postfix is itself is a multiple 8. If not, the input is rejected, otherwise, the procedure described in this paragraph is repeated. It is clear that the number of zigzags on the input is bounded by $O(\log |w|)$, and so this procedure takes at most $O(|w| \log |w|)$ steps.

In the quantum phase, the head is placed on the left end-marker, and the procedure of Figure 3 is implemented on a single qubit. It has already been shown that this procedure accepts with high probability if the number of a's that it scans corresponds to a string in L, and rejects with high probability otherwise. This procedure takes m steps. □

Moreover, POWER-EQ can be recognized by 2qcca's in linear time exactly (deterministically). The numbers i and j in each block of the form $a^i b a^j$ can be compared while going from left to right, and the overall running time is linear.

Corollary 1. *For any given* L, POWER-EQ(L) *can be recognized by 2qcca's with bounded error in linear time.*

We can obtain similar results for other computational models. For example, a one-way two-head deterministic finite automaton can easily recognize POWER-EQ in linear time. Therefore, POWER-EQ(L), for any given L, can be recognized by a 1qcfa with two input heads in linear time.

Moreover, POWER-EQ can also be recognized by a 1.5-way[6] quantum finite automaton (1.5qfa) [5] with bounded error in linear time, by a straightforward modification of techniques given in [9, 5, 15, 13]. The procedure given in Figure 3 can be implemented by such a 1.5qfa in parallel. Therefore, POWER-EQ(L), for any given L, can also be recognized by 1.5qfa's with bounded error.

[6] A 1.5qfa makes a single left-to-right pass on the input. The head is quantum and is allowed to pause on the same square for some steps, so it is possible to have superpositions of several different head positions.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de