Quantum Computing
Lecture 8

Quantum Automata and Complexity

Anuj Dawar

Models of Computation

Shor's algorithm solves, in polynomial time, a problem for which no *classical, deterministic* polynomial time algorithm is known.

What class of problems are solvable by quantum machines in polynomial time?

More generally, how does *quantum parallelism* compare with other forms of *parallelism* and *nondeterminism*.

How does the quantum model of computation affect our understanding of complexity classes?

Finite State Systems

This automaton accepts the set of strings that contain at least one b.

Its operation can be described by a pair of matrices.

$$M_a = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \qquad\qquad M_b = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}$$

$M_b M_b M_a$ describes the operation on states performed by reading the string abb.

DFAs and Matrices

Each DFA is specified by a collection of $n \times n$ matrices, where n is the number of states in the DFA, and there is one matrix for each letter.

Each column of each matrix is a unit vector with 0, 1 entries.

More generally, we can form a matrix M_w for each word w.

Multiplication of matrices corresponds to concatenation of words.

$M_w |i\rangle = |j\rangle$ if there is a path labelled w from state i to state j.

Nondeterministic Automata

This automaton accepts the same set of strings as the deterministic one.

The columns of the corresponding matrices are no longer unit vectors.

$$M_a = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$M_b = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$$

NFAs and Matrices

$$M_{bb} = \left[\begin{array}{cc} 1 & 0 \\ 2 & 1 \end{array} \right]$$

$M_w(i,j)$ gives the *number of paths* labeled w from state i to state j.

$$M_w|i\rangle = \sum_j M_w(i,j)|j\rangle$$

Probabilistic Automata

We obtain *probabilistic automata* if we allow fractional values in M_σ.
with the proviso that each column adds up to 1.

E.g. $\quad M_a = \begin{bmatrix} 0.5 & 0.0 \\ 0.5 & 1.0 \end{bmatrix}$ $\qquad M_b = \begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}$

$M_a M_b = \begin{bmatrix} 0.4 & 0.1 \\ 0.6 & 0.9 \end{bmatrix}$ gives, in position (i, j) the probability that string ba takes you from state i to state j.

Language Accepted

A probabilistic automaton $\mathcal{A}$ accepts a language L with certainty if

$$P(\mathcal{A} \text{ accepts } w) = \left\{ \begin{array}{l} 1 \text{ if } w \in L \\ 0 \text{ if } w \notin L \end{array} \right.$$

$\mathcal{A}$ accepts a language L with bounded probability if there is an $\epsilon < 1/2$ such that:

$$P(\mathcal{A} \text{ accepts } w) = \left\{ \begin{array}{ll} > 1 - \epsilon & \text{if } w \in L \\ < \epsilon & \text{if } w \notin L \end{array} \right.$$

The class of languages accepted by probabilistic automata (under either definition) is the regular languages.

Quantum Automata

Quantum finite automata are obtained by letting the matrices M_σ have complex entries.

> *We also require each of the matrices to be* unitary.

E.g. $\quad M_\sigma = \begin{bmatrix} -1 & 0 \\ 0 & i \end{bmatrix}$

M_σ is unitary since the sum of the squares of the norms in each column adds up to 1 *and* the dot product of any two columns is 0.

NB: If all matrices only have 0 or 1 entries and the matrices are unitary (i.e., the matrices are *permutation matrices*), then the automaton is *deterministic* and *reversible*.

Acceptance Probabilities

If q is the starting state of the automaton,

$$M_w |q\rangle$$

is the state after reading w. If

$$\alpha_j = \langle j | M_w | q \rangle$$

then α_j is the *probability amplitude* that the automaton reaches state j.

$|\alpha_j|^2$ is the probability that a measurement will result in state q_j.

$\sum_{j \in F} |\alpha_j|^2$ is the probability that the automaton accepts the string w.

Language Accepted

We can define language acceptance exactly or by bounded probability.

Because of the reversibility requirement, quantum automata are quite a weak model.

There are *regular* languages that cannot be accepted by a QFA.

However, QFAs may be much more *succinct* than their classical counterparts.

Interference

Consider the automaton in a one letter alphabet defined as:

$$M_a = \begin{bmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix}$$

$$M_{aa} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

While there are two distinct paths labelled *aa* from q_0 back to itself, and each has non-zero probability, the net probability of ending up in q_0 is 0.

The automaton accepts a string of odd length with probability 0.5 and a string of even length with probability 1 if its length is not a multiple of 4 and probability 0 otherwise.

Turing Machines

A *Turing machine*, in addition to the finite set of states in the automaton has an infinite *read-write* tape.

A machine is determined by an alphabet Σ, a finite set of states Q and a transition function δ which gives, for each state and symbol: a next state, a replacement symbol and a direction in which to move the tape head.

A machine has infinitely many possible *configurations* (reserving the word "state" for a member of Q).
Each configuration c is determined by a state, the contents of the tape (a finite string) and the position of the head.

Configurations and Computations

If c_0 is the configuration in the starting state, with w on the tape and the tape head at the left end of the string, w is accepted if the computation

$$c_0 \to c_1 \to \cdots \to c_f$$

eventually reaches an accepting state.

If the length of the computation is bounded by a polynomial in the length of w, the language accepted by the machine is in P.

The action of the Turing machine can equivalently be described as a linear operator M on an *infinite-dimensional* space.

The set of configurations form a basis for the space.

Nondeterministic Turing Machines

In a *nondeterministic machine*, δ determines, for each state and symbol, a *set* of next moves.

This gives rise to a tree of configurations:

A configuration may occur in several places in the tree.

The initial string w is accepted by the machine if there is some path through the tree leading to an accepting state.

If the height of the tree is bounded by a polynomial in the length of w, then the language is in NP

Probabilistic Machines

With a *probabilistic machine*, δ defines, for each current state and symbol, a *probability distribution* over the possible next moves.

The action of the machine can be defined as an infinite matrix, where the rows and columns are configurations, and each column adds up to 1.

However, how much information can be encoded in a single entry?
We require the entries α to be *feasibly computable*. That is, there is a feasibly computable f such that:

$$|f(n) - \alpha| < 2^{-n}$$

BPP

BPP is the collection of languages L for which there is a probabilistic machine M, running in polynomial time with:

$$P(M \text{ accepts } w) = \begin{cases} > \frac{2}{3} & \text{if } w \in L \\ < \frac{1}{3} & \text{if } w \notin L \end{cases}$$

The class of languages is unchanged if we replace $\frac{2}{3}$ and $\frac{1}{3}$ by $1 - \epsilon$ and ϵ, for any $\epsilon < \frac{1}{2}$, or indeed the set of all feasibly computable probabilities with $\{0, \frac{1}{2}, 1\}$.

The only inclusion relations we know are P$\subseteq$NP and P$\subseteq$BPP.

Primality testing, long known to be in NP and in BPP was shown in 2002 to be in P.

Quantum Turing Machines

With a *Quantum Turing machine*, δ associates with each state and symbol, and each possible next move, a *complex probability amplitude* (which we require to be a *feasible* complex number).

The machine can be seen as progressing through a series of stages, each of which is a superposition of configurations.

Note that the probability of c_7 occurring at time 2 may be less than the sum of the probabilities along the two paths.

We also require that the linear transformation defined by the machine is unitary.

BQP

BQP is the collection of languages L recognised by a quantum Turing machine, running in polynomial time, under the bounded probability rule.

The class BQP is not changed if we restrict the set of possible amplitudes to $\{0, \pm\frac{3}{5}, \pm\frac{4}{5}, 1\}$.

$$BPP \subseteq BQP$$

Shor's algorithm shows that the factorisation problem is in BQP. It is not known to be in BPP.

Complexity Classes

Inclusion relations among the complexity classes as we know them.