

Quantum Proofs

John Watrous

Department of Computer Science
University of Calgary

August 23, 2005

Promise problems

A **promise problem** is a computational problem where two disjoint sets of inputs (**yes** inputs and **no** inputs) must be distinguished.

For example:

GRAPH ISOMORPHISM

Input: Two simple, undirected graphs G_0 and G_1 .

Yes: G_0 and G_1 are isomorphic ($G_0 \cong G_1$).

No: G_0 and G_1 are not isomorphic ($G_0 \not\cong G_1$).

More formally: a promise problem is a pair $A = (A_{\text{yes}}, A_{\text{no}})$, with $A_{\text{yes}}, A_{\text{no}} \subseteq \{0, 1\}^*$ and $A_{\text{yes}} \cap A_{\text{no}} = \emptyset$.

We do not require $A_{\text{yes}} \cup A_{\text{no}} = \{0, 1\}^*$; there may be **don't care** inputs.

Some basic complexity classes

- P** The class of promise problems solvable in **polynomial time** on a deterministic Turing machine.
- BPP** The class of promise problems solvable in **polynomial time** on a **bounded error** probabilistic Turing machine (correct on every input with probability at least 99/100).
- PP** The class of promise problems solvable in **polynomial time** on an **unbounded error** probabilistic Turing machine (correct on every input with probability greater than 1/2).
- PSPACE** The class of promise problems solvable in **polynomial space** on a deterministic Turing machine.
- EXP** The class of promise problems solvable in **exponential time** on a deterministic Turing machine.

A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in the class NP if and only if there exists:

- a polynomial p
- a polynomial-time computable predicate V

such that these two properties are satisfied:

1. Completeness

If $x \in A_{\text{yes}}$ then there exists a string y of length $p(|x|)$ such that $V(x, y) = 1$.

The string y is a **proof** (or **certificate** or **witness**) that $x \in A_{\text{yes}}$.

2. Soundness

If $x \in A_{\text{no}}$ then $V(x, y) = 0$ for every string y of length $p(|x|)$.


Example

GRAPH ISOMORPHISM

Input: Two simple, undirected graphs G_0 and G_1 .

Yes: G_0 and G_1 are isomorphic ($G_0 \cong G_1$)

No: G_0 and G_1 are not isomorphic ($G_0 \not\cong G_1$)



Isomorphism:

$1 \rightarrow 3, 2 \rightarrow 2, 3 \rightarrow 7, 4 \rightarrow 5, 5 \rightarrow 1, 6 \rightarrow 6, 7 \rightarrow 4.$

Closure under complement?


GRAPH NON-ISOMORPHISM

Input: Two simple, undirected graphs G_0 and G_1 .

Yes: G_0 and G_1 are not isomorphic ($G_0 \not\cong G_1$).

No: G_0 and G_1 are isomorphic ($G_0 \cong G_1$).

Consider certifying that these two graphs are **non-isomorphic**:



It is not known whether or not this problem is in NP... an efficient **general** method would be required.

MA is defined similarly to NP, except that the verification procedure is probabilistic. . . a promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in MA if and only if there exists:

- a polynomial p
- a polynomial-time **probabilistic** algorithm V

such that these two properties are satisfied:


1. Completeness

If $x \in A_{\text{yes}}$ then there exists a string y of length $p(|x|)$ such that V accepts (x, y) with probability at least $99/100$.

2. Soundness

If $x \in A_{\text{no}}$ then V rejects (x, y) for every string y of length $p(|x|)$ with probability at least $99/100$.

Diagram of classes



Complete promise problems

Definition

A problem $A = (A_{\text{yes}}, A_{\text{no}})$ is **complete** for a complexity class \mathcal{C} if:

1. A is in \mathcal{C} .
2. For every other promise problem $B = (B_{\text{yes}}, B_{\text{no}})$ in \mathcal{C} , there exists a polynomial-time computable function f : such that

$$x \in B_{\text{yes}} \Rightarrow f(x) \in A_{\text{yes}},$$


$$x \in B_{\text{no}} \Rightarrow f(x) \in A_{\text{no}}.$$

Identification of one or more complete promise problems can sometimes be helpful for better understanding complexity classes.

(Not all complexity classes of promise problems have complete problems, but many do.)

Quantum circuits


In this talk we will restrict our attention to quantum circuits build from three unitary gates: **Toffoli gates**, **phase-shift gates**, and **Hadamard gates**.



This is a **universal** set of gates (and can perform reversible computations without error).

Quantum circuits

It will also be useful to consider two simple **non-unitary** gates, that correspond to **creating** and **destroying** a single qubit:




By adding these gates to those from the previous slide (Toffoli, Hadamard, and phase-shift), we get a collection that is universal in a slightly stronger sense. . .

Universal Gate Set

Any completely positive trace-preserving (i.e., **admissible**) map can be approximated with any desired accuracy.

Example: AND circuit

The following circuit takes two qubits as input and outputs one qubit:



Its operation is equivalent to measuring the input qubits in the standard basis and outputting the AND of the measurement results.


Terminology

A *type* (n, m) quantum circuit is one that takes n qubits as input and outputs m qubits.

A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in BQP if there exists a polynomial-time generated family $\{Q_x : x \in \{0, 1\}^*\}$ of type $(0, 1)$ quantum circuits such that:

1. If $x \in A_{\text{yes}}$, then $\langle 1|Q_x|1\rangle \geq 99/100$, and
2. if $x \in A_{\text{no}}$, then $\langle 1|Q_x|1\rangle \leq 1/100$.

Because each circuit Q_x takes no input, we can simply identify Q_x with a single-qubit density matrix ρ_x :



QMA: a quantum analogue of NP

A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in QMA if there exists:

- a polynomial p , and
- a polynomial-time generated family $\{Q_x : x \in \{0, 1\}^*\}$ of circuits, where each Q_x is of type $(p(|x|), 1)$,

with similar properties to the classical case:

1. Completeness

If $x \in A_{\text{yes}}$, then there exists a state ρ on $p(|x|)$ qubits such that

$$\langle 1 | Q_x(\rho) | 1 \rangle \geq \frac{99}{100}.$$


2. Soundness

If $x \in A_{\text{no}}$, then for every state ρ on $p(|x|)$ qubits it holds that

$$\langle 1 | Q_x(\rho) | 1 \rangle \leq \frac{1}{100}.$$

QMA: a quantum analogue of NP

The circuit family $\{Q_x : x \in \{0, 1\}^*\}$ is a **verification procedure**.



Completeness: if $x \in A_{\text{yes}}$ then there exists a **quantum proof** ρ of this fact; $\langle 1|Q_x(\rho)|1\rangle \geq 99/100$.

Soundness: if $x \in A_{\text{no}}$ then **every** state ρ fails to prove otherwise; $\langle 1|Q_x(\rho)|1\rangle \leq 1/100$.

A more refined definition of QMA

Given a polynomial p (the **proof length**) and polynomial-time computable functions a and b (the **completeness and soundness probabilities**), let us define

$$\text{QMA}_p(a, b)$$

to be the class of all problems $A = (A_{\text{yes}}, A_{\text{no}})$ such that there exists a polynomial-time generated family $\{Q_x : x \in \{0, 1\}^*\}$ of circuits, where each Q_x is of type $(p(|x|), 1)$, such that:

1. Completeness

If $x \in A_{\text{yes}}$, then $\langle 1|Q_x(\rho)|1\rangle \geq a(|x|)$ for some $p(|x|)$ -qubit state ρ .

2. Soundness

If $x \in A_{\text{no}}$, then $\langle 1|Q_x(\rho)|1\rangle \leq b(|x|)$ for every $p(|x|)$ -qubit state ρ .

Some basic facts about QMA

1. Strong error reduction.

For any choice of the proof length p and completeness and soundness probabilities a and b with

$$a(n) - b(n) \geq \frac{1}{q(n)}$$

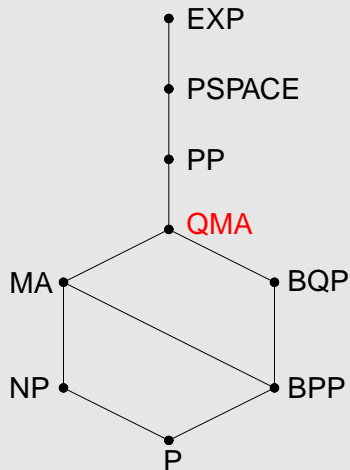
for some polynomial q , it holds that

$$\text{QMA}_p(a, b) = \text{QMA}_p(1 - 2^{-r}, 2^{-r})$$

for every polynomial r . [MARRIOTT & W., 2004.]

2. $\text{QMA} \subseteq \text{PP}$. [KITAEV & W., 2000; VYALYI, 2003; MARRIOTT & W., 2004]

Diagram of classes



2-LOCAL HAMILTONIAN

Input: A description of a 2-local Hamiltonian H on n qubits and rational numbers a, b with $b - a \geq 1/\text{poly}$.

Yes: $\lambda_{\min}(H) \leq a$.

No: $\lambda_{\min}(H) \geq b$.

Theorem [KEMPE, KITAEV & REGEV, 2004.]

2-LOCAL HAMILTONIAN is complete for QMA.

Group-theoretic problems

Let G be a **finite group** whose elements can be represented (uniquely) by strings of a given length n .

Efficient computation of group operations:

Given two elements $g, h \in G$, it is assumed that the group operations can be efficiently implemented by quantum circuits:

1. **Multiplication:** $|g\rangle |h\rangle \mapsto |g\rangle |gh\rangle$.
2. **Inverse:** $|g\rangle \mapsto |g^{-1}\rangle$.

Abstraction:

It is sometimes helpful to view such a group as a **black box group**; the group operations are performed by a black box (or group oracle), and string representatives of elements are independent of group structure.

GROUP MEMBERSHIP

Input: Group elements g_1, \dots, g_k and h of G .

Yes: $h \in \langle g_1, \dots, g_k \rangle$.

No: $h \notin \langle g_1, \dots, g_k \rangle$.

- GROUP MEMBERSHIP \in NP [BABAI AND SZEMERÉDY, 1984]
The proof follows from the *Reachability Lemma*: every element in the subgroup $\langle g_1, \dots, g_k \rangle$ has a short *straight-line program* that starts with g_1, \dots, g_k .
- GROUP NON-MEMBERSHIP is **not** known to be in NP (or in MA).
There are group oracles relative to which it is provably not the case [BABAI, 1991; W., 2000].

Quantum proofs for non-membership

Theorem [W., 2000]

GROUP NON-MEMBERSHIP is in QMA.

The idea behind the proof of this theorem is simple—the quantum state that proves

$$h \notin H \stackrel{\text{def}}{=} \langle g_1, \dots, g_k \rangle$$


will be the uniform pure state over the elements of H :

$$|H\rangle = \frac{1}{\sqrt{|H|}} \sum_{a \in H} |a\rangle.$$

(It is independent of the element h .)

Quantum proofs for non-membership

Suppose that you have a copy of the state $|H\rangle$. You can use this state to efficiently test membership of h in H as follows. . .




Case 1: $h \in H$. We have

$$M_h |H\rangle = |hH\rangle = |H\rangle;$$

the controlled-multiplication has **no effect**. As $H^2 |0\rangle = |0\rangle$, so the measurement outcome is **0** (with certainty).

Quantum proofs for non-membership

Suppose that you have a copy of the state $|H\rangle$. You can use this state to efficiently test membership of h in H as follows. . .



Case 2: $h \notin H$. We have

$$M_h |H\rangle = |hH\rangle \perp |H\rangle ;$$

the controlled-multiplication **acts as a measurement** of B. Both before and after the second Hadamard transform, the register B is therefore **totally mixed**. The measurement outcome is a **uniform random bit**.

But we can't trust the proof. . .


Suppose that $|\psi\rangle$ is the quantum state that supposedly proves $h \notin H$. Unfortunately we cannot trust that $|\psi\rangle = |H\rangle$, so we need to process $|\psi\rangle$ before running the membership test.

Imagine that instead of running the membership test with h , we run the test with some element $a \in H$. It should reveal that $a \in H$, because it is!

If the test indicates $a \notin H$, then we know $|\psi\rangle \neq |H\rangle$; **the proof is invalid so reject.**

Conditioned on the test indicating $a \in H$, what happens to the proof?

Modified proof



Repeat for a well-chosen set of elements a_1, \dots, a_k ; conditioned on success for each test, we will have a state very close to

$$\sum_{a \in H} M_a |\psi\rangle \quad (\text{normalized}).$$

This state is invariant under left multiplication by elements in H ; if $h \in H$, the test will falsely conclude $h \notin H$ with very small probability.

Interactive proof systems

An **interactive proof system** involves an interaction between two parties, the **prover** and the **verifier**, concerning an input string x to a promise problem $A = (A_{\text{yes}}, A_{\text{no}})$.

The Prover

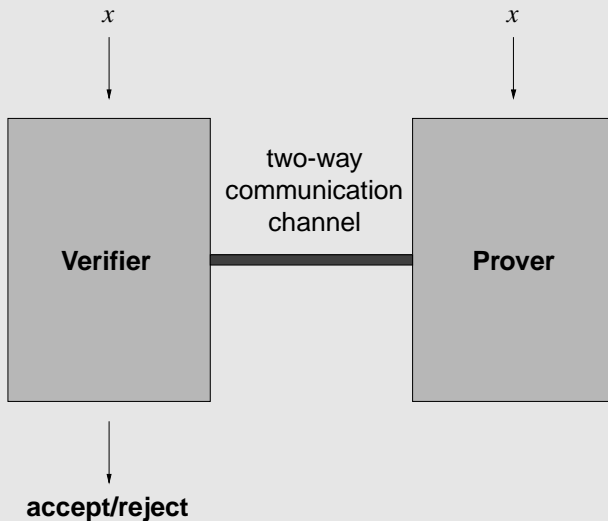
- Goal is to convince the verifier that $x \in A_{\text{yes}}$.
- The prover is **not trustworthy**—it will try to prove $x \in A_{\text{yes}}$ even when this is not the case.
- The prover is **computationally unbounded**.

The Verifier

- Goal is to check the validity of the prover's argument that $x \in A_{\text{yes}}$.
- The verifier is **computationally bounded**. It essentially has the power of BPP (classical) or BQP (quantum).

[GOLDWASSER, MICALI & RACKOFF, 1985; BABAI, 1985]

Informal picture of an interactive proof



A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ **has an interactive proof system** if there exists a verifier V that satisfies similar conditions to the non-interactive case:

1. Completeness.

If $x \in A_{\text{yes}}$, then there is a prover P that convinces V to accept x (with high probability).

2. Soundness.

If $x \in A_{\text{no}}$, then no prover P can convince V to accept x (except with small probability).

Informal example

Claim

Coke and Pepsi taste different.

Suppose you believe this claim, and indeed can taste the difference. (Substitute “salt and sugar” if it helps with the example.)

How could you convince a skeptic?

Non-interactive proof: hopeless.

Interactive proof: easy. Let the skeptic run a **blind taste test**; when you win every time, he should be convinced.

Similar protocol for graph non-isomorphism

GRAPH NON-ISOMORPHISM

Input: Two simple, undirected graphs G_0 and G_1 .

Yes: G_0 and G_1 are not isomorphic ($G_0 \not\cong G_1$).

No: G_0 and G_1 are isomorphic ($G_0 \cong G_1$).

There is a simple (classical) interactive proof requiring just one question and response. [GOLDREICH, MICALI & WIGDERSON, 1991]

1. The verifier randomly chooses a bit $b \in \{0, 1\}$ and a permutation $\sigma \in S_n$, and lets $H = \sigma(G_b)$.
2. The verifier sends H to the prover, and challenges him to guess whether $b = 0$ or $b = 1$. If the prover guesses correctly, the verifier **accepts**, otherwise he **rejects**.

(Repeat many times *in parallel* to decrease the error probability.)

Facts about interactive proofs

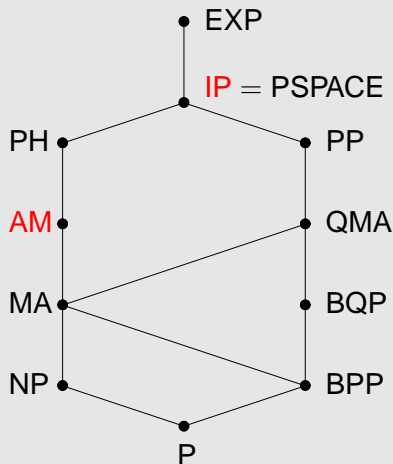
- Complexity classes based on classical interactive proofs are robust with respect to the choice of completeness and soundness probabilities.
- The complexity class containing all promise problems having classical interactive proof systems is denoted IP . It is known that $IP = PSPACE$. [FELDMAN, 1986; LUND, FORTNOW, KARLOFF & NISAN, 1990; SHAMIR, 1990]
- Denote the class of promise problems having classical interactive proofs with at most m messages by $IP(m)$. For any **constant** $m \geq 2$ it holds that

$$IP(m) = IP(2) = AM \subseteq PH.$$

[BABAI, 1985; GOLDWASSER & SIPSER, 1989]

(Known classical protocols for $PSPACE$ require a polynomial number of messages.)

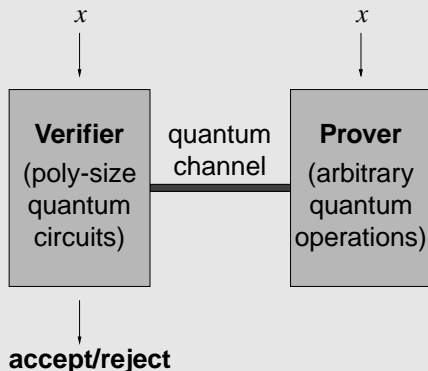
Diagram of complexity classes



Quantum interactive proof systems

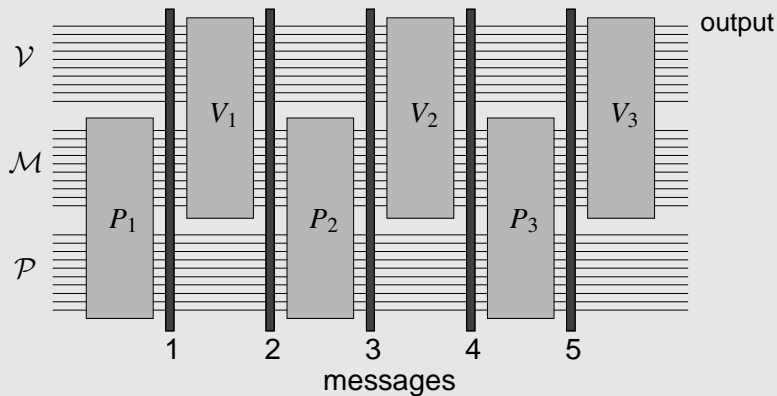
The **quantum interactive proof system** model works exactly the same as the classical model, except that the prover and verifier may **exchange and process quantum information**.

General assumptions and notions of completeness and soundness are unchanged. . .



Circuit for quantum interactive proof

The model is formalized in terms of quantum circuits. . .



Facts about quantum interactive proofs


- Complexity classes based on quantum interactive proofs are robust with respect to the choice of completeness and soundness probabilities (similar to classical case, but harder to prove).
- Let $\text{QIP}(m)$ denote the class of promise problems having m message quantum interactive proof systems. Then

$$\text{QIP} \stackrel{\text{def}}{=} \text{QIP}(\text{poly}) = \text{QIP}(3).$$

[KITAEV & W., 2000]

- It holds that $\text{PSPACE} \subseteq \text{QIP} \subseteq \text{EXP}$. [KITAEV & W., 2000]
- We have already discussed $\text{QIP}(1) = \text{QMA}$, and not too much is known about $\text{QIP}(2)$ except trivialities. . .

Diagram of complexity classes



A simple complete problem for QIP

Given two admissible maps Q and R (with the same output dimension), define their **max-fidelity** as

$$F_{\max}(Q, R) = \max_{\rho, \xi} F(Q(\rho), R(\xi))$$

(maximum is over density matrices ρ and ξ of appropriate dimensions).

CLOSE IMAGES


Input: Quantum circuits Q and R (of the same type).

Yes: $F_{\max}(Q, R) = 1$.

No: $F_{\max}(Q, R) \leq 1/10$.

Completeness of close images

Given a 3-message quantum interactive proof for some promise problem:



Circuits implementing Q_0 and Q_1 are easy to obtain given a description of V_1 and V_2 , and contain all the information we need:

$$\max_{P_1, P_2} \text{Prob}[V \text{ accepts}] = F_{\max}(Q_0, Q_1)^2$$

Trace norm and distinguishability

The **trace norm** of a matrix X is the sum of the singular values of X . Equivalently,

$$\|X\|_{\text{tr}} = \text{tr} \sqrt{X^\dagger X}.$$

This norm relates closely to the probability with which two mixed states ρ_0 and ρ_1 can be distinguished. . .

Distinguishability of states

Suppose $\{E_0, E_1\}$ is any binary-valued measurement and $\xi \in \{\rho_0, \rho_1\}$ is chosen uniformly. The probability the measurement correctly identifies ξ is at most

$$\frac{1}{2} + \frac{1}{4} \|\rho_0 - \rho_1\|_{\text{tr}}.$$

This bound is achieved for an optimal measurement.

Circuit distinguishability

Suppose we have two admissible maps Q_0 and Q_1 , both of type (n, m) .

Is there a similar notion of distance between Q_0 and Q_1 to the trace distance between mixed states ρ_1 and ρ_2 ?

One possibility: let the trace distance induce a norm on super-operators:

$$\|Q_0 - Q_1\|_{\text{tr}} = \max_{\xi} \|Q_0(\xi) - Q_1(\xi)\|_{\text{tr}}.$$

(Maximum over density matrices on n qubits.)

Bad choice...

So close and yet so far...

Let Q_0 and Q_1 be mappings of type (n, n) defined as follows:

$$Q_0(X) = \frac{1}{2^n + 1} ((\text{tr } X)I + X^T), \quad Q_1(X) = \frac{1}{2^n - 1} ((\text{tr } X)I - X^T).$$

These are both admissible maps (and could be implemented by circuits).

- For every mixed state ξ it holds that $\|Q_0(\xi) - Q_1(\xi)\|_{\text{tr}} \leq \frac{4}{2^n + 1}$.
- Let



$$|\psi\rangle = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle |i\rangle.$$

Then

$$\|(Q_0 \otimes I)(|\psi\rangle \langle \psi|) - (Q_1 \otimes I)(|\psi\rangle \langle \psi|)\|_{\text{tr}} = 2;$$


$(Q_0 \otimes I)(|\psi\rangle \langle \psi|)$ and $(Q_1 \otimes I)(|\psi\rangle \langle \psi|)$ are **perfectly distinguishable**.

So close and yet so far...



$\rho_0 \approx \rho_1$ (for any choice of ξ)

So close and yet so far...



$\rho_0 \perp \rho_1$ (i.e., they are perfectly distinguishable)

Kitaev's "diamond" norm

Given two admissible maps Q_0 and Q_1 , both of type (n, m) , define

$$\begin{aligned}\|Q_0 - Q_1\|_{\diamond} &= \|Q_0 \otimes I_n - Q_1 \otimes I_n\|_{\text{tr}} \\ &= \max_{\xi} \|(Q_0 \otimes I_n)(\xi) - (Q_1 \otimes I_n)(\xi)\|_{\text{tr}}\end{aligned}$$

where I_n is the identity super-operator on n qubits, and the maximum is over all density operators ξ on $2n$ qubits.

This norm is called the **diamond norm**, and it has several remarkable properties. . .

The **diamond-distance** between admissible maps may be viewed as the sensible analogue of the trace-distance between density operators.

Quantum circuit distinguishability

Consider the following problem (where ε is a small positive constant).

QUANTUM CIRCUIT DISTINGUISHABILITY

Input: Quantum circuits Q_0 and Q_1 (of the same type).

Yes: $\|Q_0 - Q_1\|_{\diamond} \geq 2 - \varepsilon$.

No: $\|Q_0 - Q_1\|_{\diamond} \leq \varepsilon$.

Informally, this problem asks whether two physical processes (described by quantum circuits) act nearly identically or not.

Theorem [ROSGEN & W., 2004]

QUANTUM CIRCUIT DISTINGUISHABILITY is complete for QIP.

Open problems

1. Place interesting problems in QMA.
 - Is GRAPH NON-ISOMORPHISM in QMA?
 - Is GROUP ORDER in QMA?
2. Many questions about the classes QMA, QIP(2), and QIP remain unanswered.
 - Is $\text{QIP}(2) \subseteq \text{PSPACE}$?
 - Improve $\text{PSPACE} \subseteq \text{QIP} \subseteq \text{EXP}$.
 - Is QIP closed under complementation?
3. There are several interesting variants of these models for which little is known:
 - “Multiple Merlins”... are two quantum proofs better than one?
 - Irrefutable quantum arguments...
 - Multiprover interactive proofs...
 - Zero-knowledge interactive proofs...