Hard Quantum Problems

Bill Rosgen

Centre for Quantum Technologies
National University of Singapore

February 26, 2013

Overview

- As the size of the quantum systems we can control grows, it is increasingly difficult to characterize them: the number of parameters grows exponentially in the number of qubits.
- Complexity theory can be used to identify problems we cannot solve in time polynomial in the number of qubits.
- Worst-case results are not impossibility proofs, but indications that additional structure must be exploited

Outline

Computational problems

A *problem* (or *language*) P is a pair of sets P_{yes} and P_{no} such that

1. $P_{\text{yes}}, P_{\text{no}} \subseteq \{0,1\}^*$
2. $P_{\text{yes}} \cap P_{\text{no}} = \emptyset$

An *algorithm* A is some process that takes as input $x \in \{0,1\}^*$ and either accepts or rejects. A solves a problem P if

1. If $x \in P_{\text{yes}}$ then $\Pr[A(x) \text{ accepts}] \geq 1 - \varepsilon$,
2. If $x \in P_{\text{no}}$ then $\Pr[A(x) \text{ accepts}] \leq \varepsilon$.

We measure the running time of A as a function of $|x|$.

- These is a *worst-case* definition. We require an algorithm to work correctly for *every* $x \in P_{\text{yes}} \cup P_{\text{no}}$

[1] strictly: this defines the *promise problems*, but I'll ignore this distinction.

Complexity classes

A complexity class is the set of all problems solved by algorithms in some model of computation.

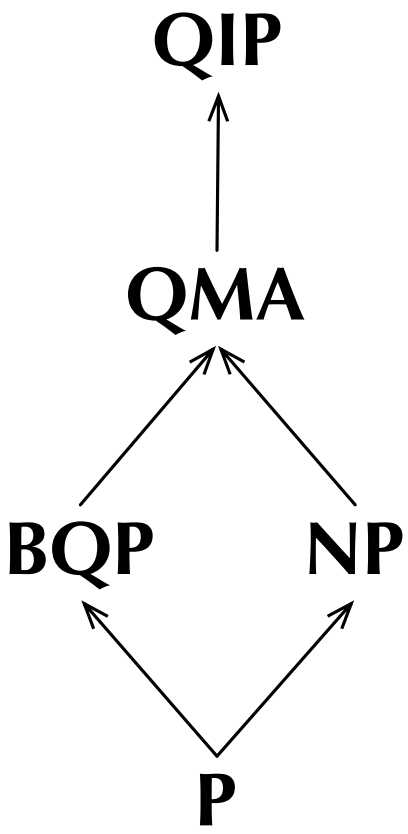

BQP: Efficient quantum computation

QMA: Efficient quantum computation with access to an untrusted quantum proof

QIP: Efficient quantum computation with interactive access to untrusted party (equal to **PSPACE** [JJUW10])

(*Efficient* always means polynomial time in the input length.)

Reductions

We need a tool to compare the hardness of different problems.

A problem P reduces to a problem Q, written $P \leq Q$ if there is a an efficient transformation $f : \{0,1\}^* \to \{0,1\}^*$ such that

1. If $x \in P_{\text{yes}}$ then $f(x) \in Q_{\text{yes}}$.
2. If $x \in P_{\text{no}}$ then $f(x) \in Q_{\text{no}}$.

Intuitively: P is no harder than Q, since we can solve P using only f and an algorithm for Q.

Two problems P, Q are *equivalent* if $P \leq Q$ and $Q \leq P$.

Complete problems

Most complexity classes have complete problems.

P is complete for a class $\mathbf{C}$ if

1. $Q \leq P$ for all $Q \in \mathbf{C}$
2. $P \in \mathbf{C}$.

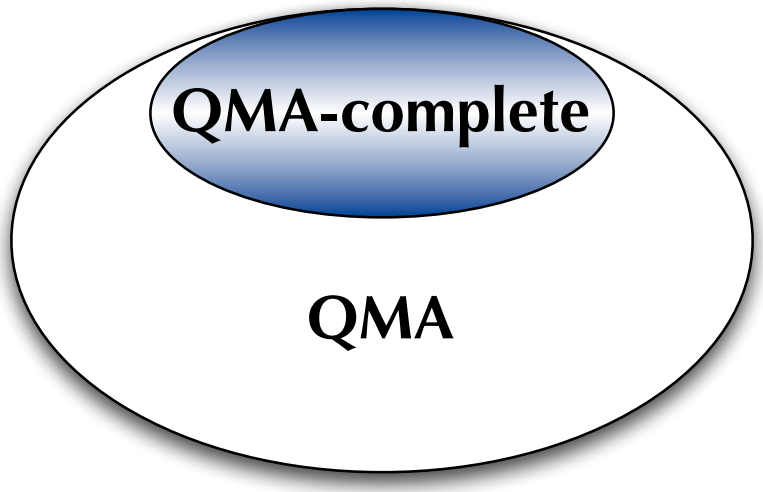

- A proof of completeness shows that a problem is 'hard'
- Complete problems characterize computational models.
- Remember: these are *worst-case* results.

How is a hardness result useful?

Suppose you want to decide if some local constraints can be satisfied by some global quantum state:

- In general this is **QMA**-complete [Liu06], and so there is no general algorithm[1]
- This does *not* mean that every instance is hard, only that there exist hard instances
- This *does* imply that you will need to use some additional structure or information

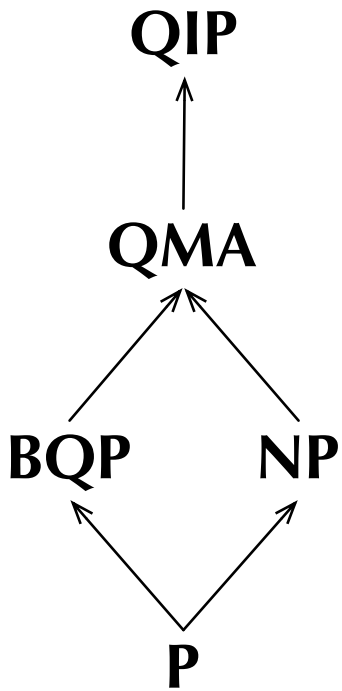

[1] Given the usual complexity-theoretic assumptions ...

Quantum (mixed-state) circuits

Often the input x to a problem will be a state or a channel. As a binary representation we will use the circuit model.

- 'Efficient' states and channels have circuits of size $\log d$, but density matrices and Kraus operators are always large.
- Given a circuit we can run it in quantum polynomial time.

We use any standard set of unitary gates plus three gates:

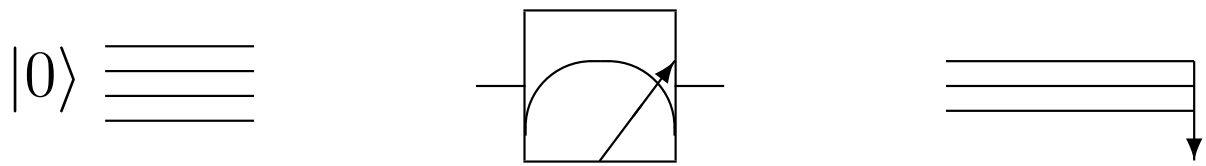

This model allows us to represent *any* quantum channel.

Circuits are generic

- The fact that we use circuits does not limit hardness results
- Given a problem in the circuit model, we can reduce it to any other model that can simulate circuits:

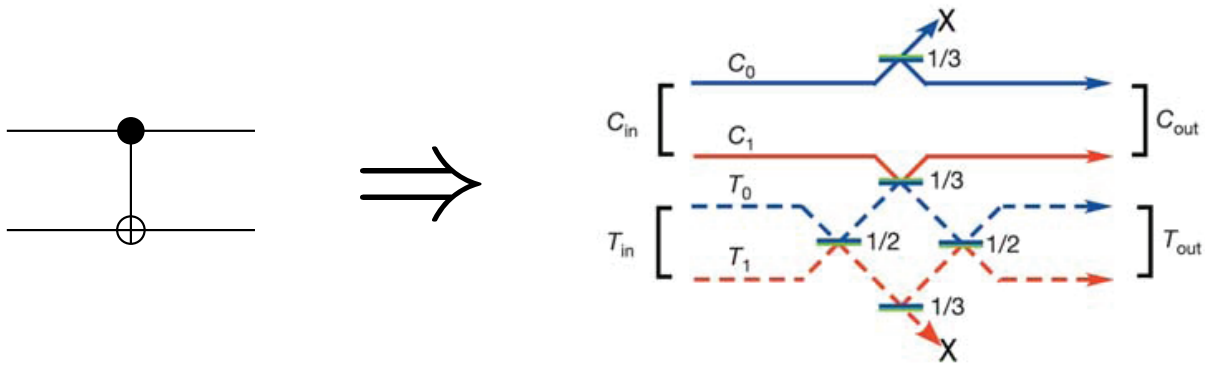

- Circuit problems no harder than problems in any model that can simulate circuits.
- Circuits give succinct descriptions of states and channels
- These descriptions can be implemented (in theory)

[1] Gate implementation [O'Brien et al., Nature 2003]

Outline

QMA

QMA: problems that can be *verified* with a quantum computer.

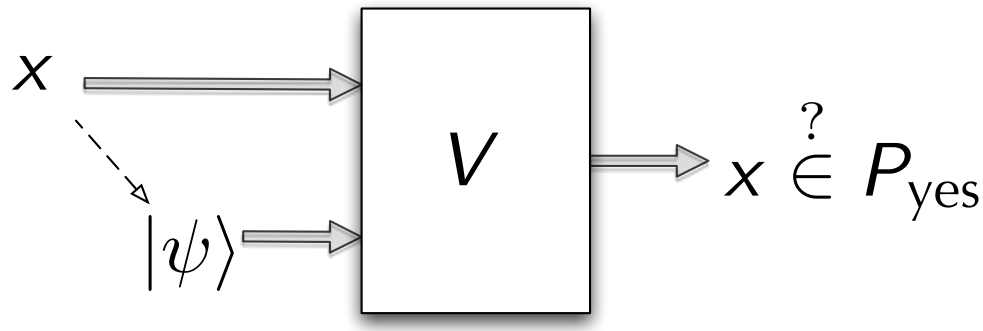

To decide if $x \in P_{\text{yes}}$, the Verifier (*Arthur*) receives a proof $|\psi\rangle$ from the Prover (*Merlin*) and runs some efficient circuit V.

- $|\psi\rangle$ may depend on x.
- Merlin is not computationally bounded.
- Merlin is not trustworthy.

QMA (definition)

QMA is the set of all problems P for which there exist polynomial-time unitary circuits V such that:

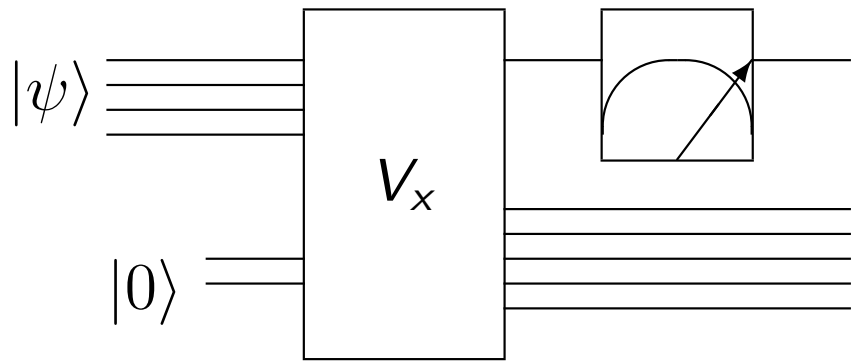

- When $x \in P_{\text{yes}}$, $\exists|\psi\rangle$ such that outcome is $|1\rangle$ with Pr $\geq 1-\varepsilon$
- When $x \in P_{\text{no}}$, $\forall|\psi\rangle$ the outcome is $|1\rangle$ with Pr $\leq \varepsilon$

QMA as a problem

This definition can also be viewed as a complete problem:

Problem (QMA)

On input C a circuit with one output qubit, decide between:

- *Yes: there exists $|\psi\rangle$ such that $\| C(|\psi\rangle) - |1\rangle\langle 1| \|_{\text{tr}} \leq \varepsilon$.*
- *No: for any $|\psi\rangle$, $\| C(|\psi\rangle) - |0\rangle\langle 0| \|_{\text{tr}} \leq \varepsilon$.*

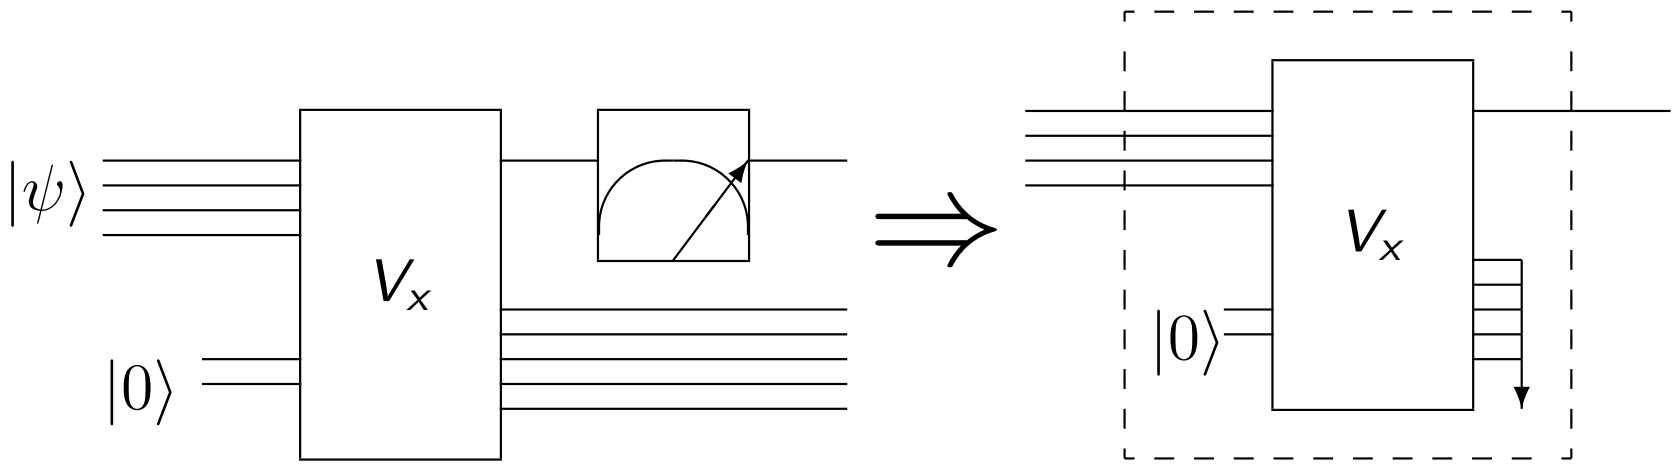

Complexity of Local Hamiltonians

Problem (k-local Hamiltonian (KSV02))

On input $a < b \in \mathbb{R}$ and k-local $H = \sum_{i=1}^{r} H_i$ on n qubits, with $r \in \text{poly}(n)$ and $\|H_i\| \leq \text{poly}(n)$ for each i, decide between:

- *Yes: smallest eigenvalue of H is less than a*
- *No: smallest eigenvalue of H is more than b*

Results:

- 5-local Hamiltonian is **QMA**-complete [KSV02]
- 2-local Hamiltonian is **QMA**-complete [KKR06]
- 2-local Hamiltonian with only nearest-neighbour interactions on a 2D grid is **QMA**-complete [OT05]
- Nearest-neighbour interactions on a 1D grid is **QMA**-complete, but uses local dimension 12 [AGIK07]

Identity Testing

Given two *unitary* circuits, are they (approximately) the same?

Equivalent to testing if a circuit is close to the identity, since

$$\| U - V \| = \| V^* U - \mathbb{1} \|$$

Problem (Non-identity check (JWB05))

Given unitary U as a circuit, decide between

- *Yes: there exists $|\psi\rangle$ such that*
$$\| U|\psi\rangle\langle\psi|U^* - |\psi\rangle\langle\psi| \|_{\text{tr}} \geq 2 - \varepsilon$$
- *No: for all $|\psi\rangle$*
$$\| U|\psi\rangle\langle\psi|U^* - |\psi\rangle\langle\psi| \|_{\text{tr}} \leq \varepsilon$$

This is **QMA**-complete [JWB05].

(Non-)Isometry Testing

How hard is it to decide if a channel Φ is close to unitary?

Problem (Non-isometry)

Given a circuit Φ, decide between:

- *Yes: For some $|\psi\rangle$, $(\Phi \otimes I)(|\psi\rangle\langle\psi|)$ is $(1-\varepsilon)$ far from pure,*
- *No: All pure-state inputs to $\Phi \otimes I$ result in ε-pure outputs.*

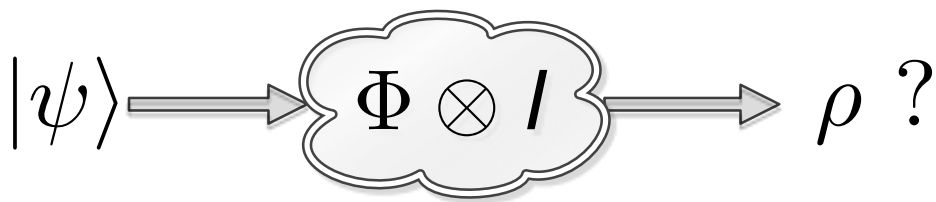

- This is **QMA**-complete [R11].
- Average case is easy:
  - find the purity of $\Phi \otimes I$ on a maximally entangled state

Hardness of Non-isometry Testing

- Argue by reduction from an arbitrary problem in **QMA**
- Add to the verifier a controlled depolarizing channel:

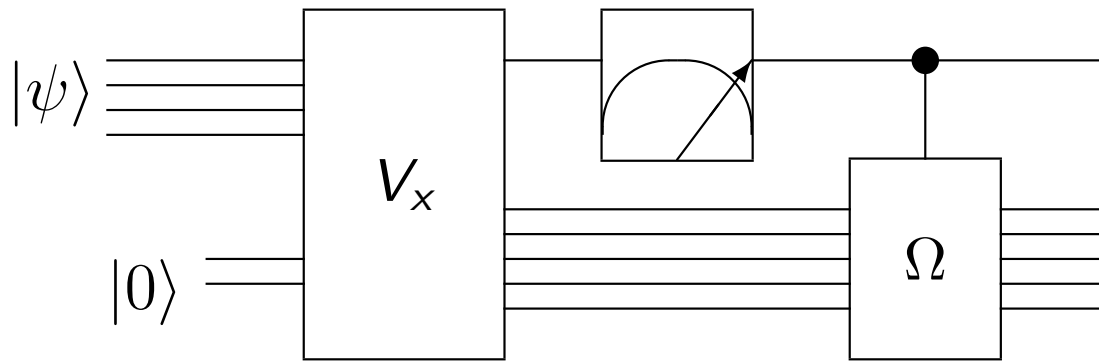

- if $x \in QMA_{\text{yes}}$, there is $|\psi\rangle$ with output $(1-2\varepsilon)$-far from pure
- if $x \in QMA_{\text{no}}$, output state is 2ε-pure for any $|\psi\rangle$

Non-isometry Testing in QMA

How to put the problem into **QMA**?
Given two *unentangled* copies of the same state $|\psi\rangle$:

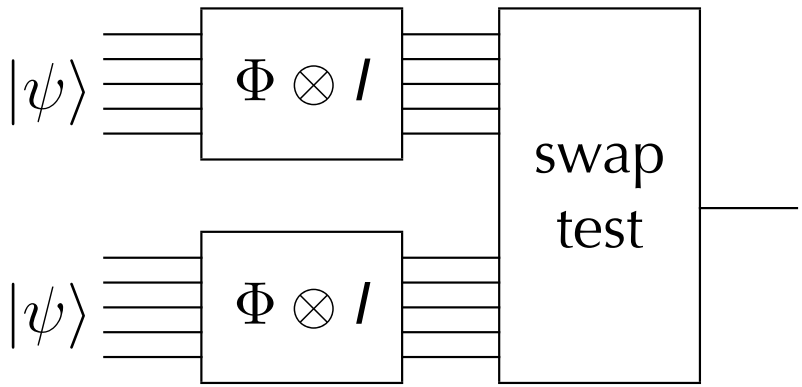

The output is 'antisymmetric' with probability

$$\frac{1}{2} - \frac{1}{2}\,\mathrm{tr}\left[(\Phi \otimes I)(|\psi\rangle\langle\psi|)^2\right]$$

Proof requires a more complicated technique to deal with entangled states, but the argument is quite technical.

Detecting Insecure Encryption

How hard is it to check that an encryption system is secure?

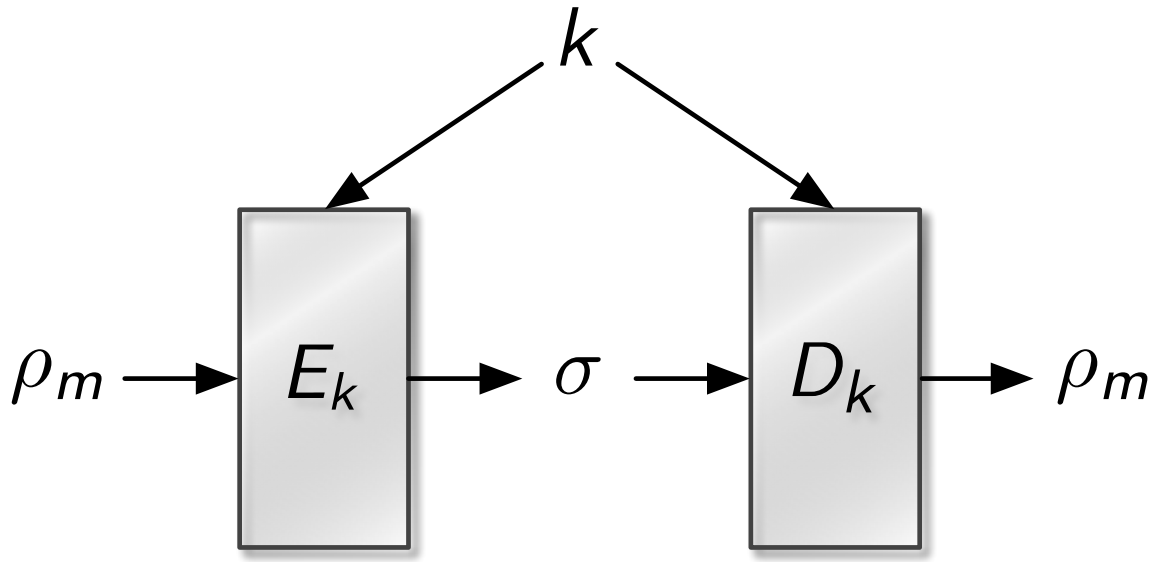

Given E_k we want to check that for any ρ_m and a random key $k \in \{0,1\}^r$, σ has almost no information about ρ_m.

More notation ...

Formalizing this requires the *diamond norm* on channels:

$$\|\Phi\|_\diamond = \|\Phi \otimes I\|_{\mathrm{tr}} = \max_{|\psi\rangle} \|(\Phi \otimes I)(|\psi\rangle\langle\psi|)\|_{\mathrm{tr}}$$

Given one of Φ_1 and Φ_2 as a black box, the probability of identifying the channel with a single use is:

$$\frac{1}{2} + \frac{\|\Phi_1 - \Phi_2\|_\diamond}{4}.$$

Detecting Insecure Encryption

Problem (Detecting Insecure Encryption)

Given a circuit E that takes a quantum input (of dimension d) and a secret key $k \in \{1, \ldots, 2^r\}$, decide between

- *Yes: There exists a subspace S of dimension $\sqrt{d}$ such that for all $|\psi\rangle \in S$ and all k*

$$\| E_k(|\psi\rangle\langle\psi|) - |\psi\rangle\langle\psi| \otimes |0\rangle\langle 0| \|_{\mathrm{tr}} \leq \varepsilon$$

- *No:*

$$\left\| \frac{1}{2^r} \sum_k E_k - \Omega \right\|_\diamond \leq \varepsilon$$

Either the encryption is almost perfect, or it fails completely on a large subspace. This problem is **QMA**-complete [R12].

Testing Quantum Circuits

Given a circuit C, does it behave like C_1 on a large subspace of the input, or does it behave like C_2 everywhere?

Problem (Circuit Testing (CT))

Let C_0 and C_1 be circuits of the same size as the input circuit C

- *Yes: There exists a subspace S of dimension $d^{(1-\delta)}$ such that for all ρ on $S \otimes \mathcal{R}$*

$$\|(C \otimes I_{\mathcal{R}} - C_1 \otimes I_{\mathcal{R}})(\rho)\|_{\mathrm{tr}} \leq \varepsilon$$

- *No: $\|C - C_2\|_{\diamond} \leq \varepsilon$, i.e. for any ρ*

$$\|(C \otimes I_{\mathcal{R}} - C_2 \otimes I_{\mathcal{R}})(\rho)\|_{\mathrm{tr}} \leq \varepsilon$$

This problem is **QMA**-hard for all families of circuits C_1, C_2 for which $\mathrm{CT}_{\mathrm{yes}} \cap \mathrm{CT}_{\mathrm{no}} = \emptyset$ [R12].

Outline

Quantum interactive proof systems

QIP: problems that can be *interactively verified*.

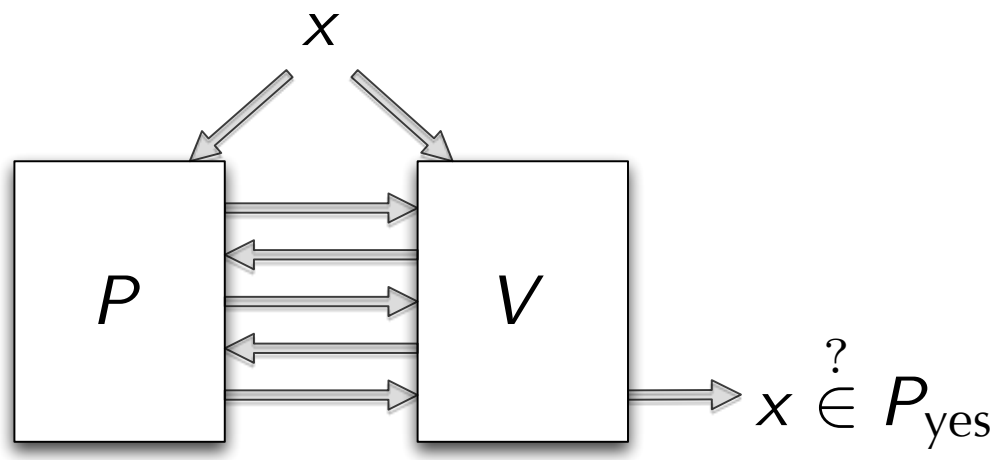

To decide if $x \in P_{\text{yes}}$, the Verifier communicates with a Prover while performing some efficient quantum computation.

- Both parties know x.
- Prover is not computationally bounded (or trustworthy).
- Both parties have private memory.
- Equal to **PSPACE** [JJUW10]

Close Images

Problem (Close Images (KW00))

Given two circuits Φ_1, Φ_2, decide between

- *Yes: There exist ρ, σ such that $\mathsf{F}(\Phi_1(\rho), \Phi_2(\sigma)) \geq 1 - \epsilon$.*
- *No: For all ρ, σ, $\mathsf{F}(\Phi_1(\rho), \Phi_2(\sigma)) \leq \epsilon$.*

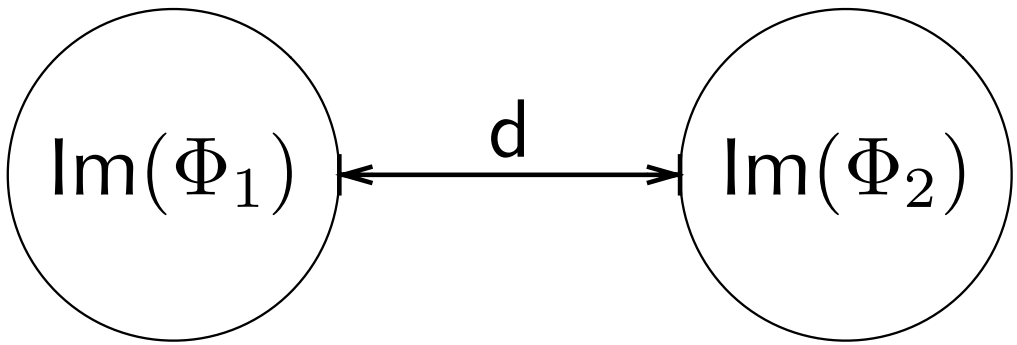

- The **QIP**-completeness of this problem follows from the fact that **QIP** can be reduced to 3 messages [KW00].
- Problem is still hard for log-depth circuits [R08].

Classical and quantum circuit distinguishability

- What is the complexity of distinguishing two circuits?
  - i.e. deciding if they have inputs on which they disagree.
  - i.e. determining if $\|\Phi_1 - \Phi_2\|_\diamond$ is large or small.

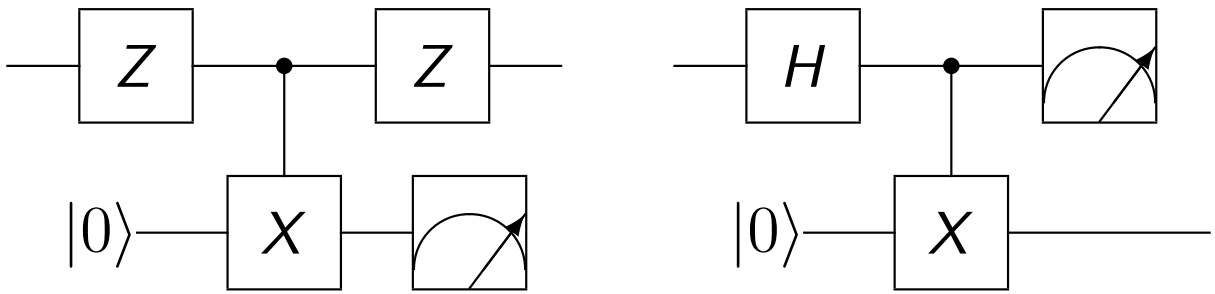

- Classical circuits: **NP**-complete
- Unitary circuits: **QMA**-complete [JWB05]
- For general channels this is **QIP**-complete [RW05]

Distinguishing quantum circuits

Problem (Quantum Circuit Distinguishability)

Given two circuits Φ_1, Φ_2, decide between

- *Yes:* $\|\Phi_1 - \Phi_2\|_\diamond \geq 2 - \epsilon$,
- *No:* $\|\Phi_1 - \Phi_2\|_\diamond \leq \epsilon$.

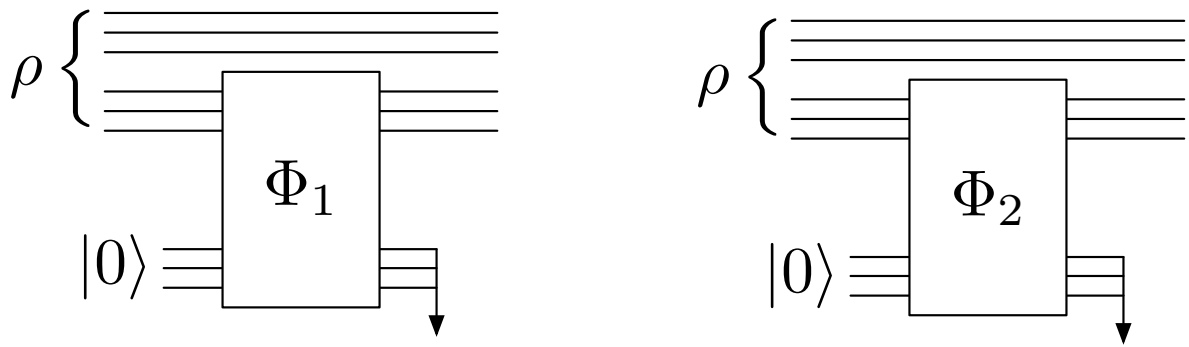

- Less formally: is there and input ρ on which $\Phi_1 \otimes I$ and $\Phi_2 \otimes I$ produce (almost) orthogonal outputs?
- This is hard for **QIP** [RW05], even for log-depth circuits [R08], and can also be used for bit-commitment [CKR11].

Random-unitary channels

Distinguishability remains **QIP**-complete when restricted to convex mixtures of unitary channels [R08].

- Φ is mixed-unitary if $\Phi(\rho) = \sum_i p_i U_i \rho U_i^*$, where the p_i form a probability distribution.

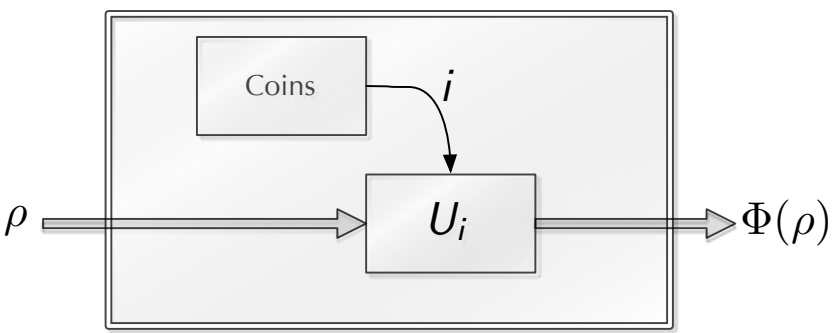

Examples:

- Depolarizing channel: $\frac{1}{4}(\rho + X\rho X + Z\rho Z + XZ\rho ZX)$
- Dephasing channel: $\frac{1}{2}(\rho + Z\rho Z)$

Degradable channels

- A channel is degradable if there is a channel that takes the output state to environment state.

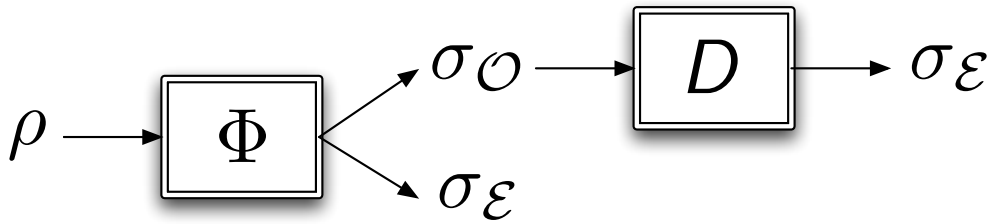

- For any channel Φ, construct the degradable channel[1]

$$\Phi'(\rho) = \frac{1}{2}|0\rangle\langle 0| \otimes \rho + \frac{1}{2}|1\rangle\langle 1| \otimes \Phi(\rho)$$

- This satisfies $\|\Phi' - \Psi'\|_\diamond = \frac{1}{2}\|\Phi - \Psi\|_\diamond$
- After some amplification, this shows that distinguishing degradable channels is **QIP**-complete [R10].
- A similar construction applies to antidegradable channels.

[1] This construction can be found in [CRS08]

Open problems / discussion

- This talk has listed many of the problems that we know to be hard in a world with large-scale quantum computers.
  - Classically we know *thousands* of **NP**-complete problems.
  - Much more interesting is understanding interesting or practical problems, i.e. what is the simplest form of tomography that is hard?
- Classically we have approximation algorithms for many hard problems. This is still largely open in the quantum case.
- These results are all worst-case hardness proofs. Is there anything useful we can prove about average-case hardness?