Testing product states and managing multiple Merlins

Ashley Montanaro

Centre for Quantum Information and Foundations,
University of Cambridge, UK

Talk based on joint work with Aram Harrow

EPSRC
Engineering and Physical Sciences Research Council

arXiv:1001.0017v3, FOCS'10

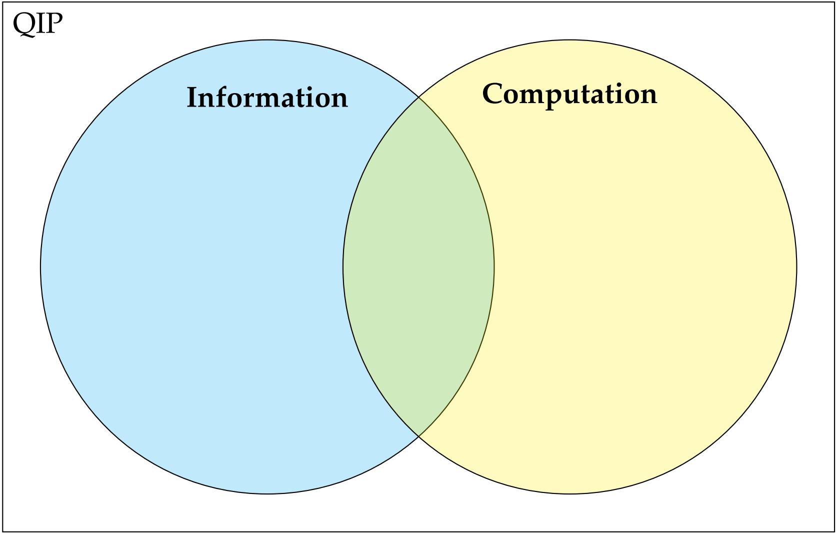
QIP
Information
Computation

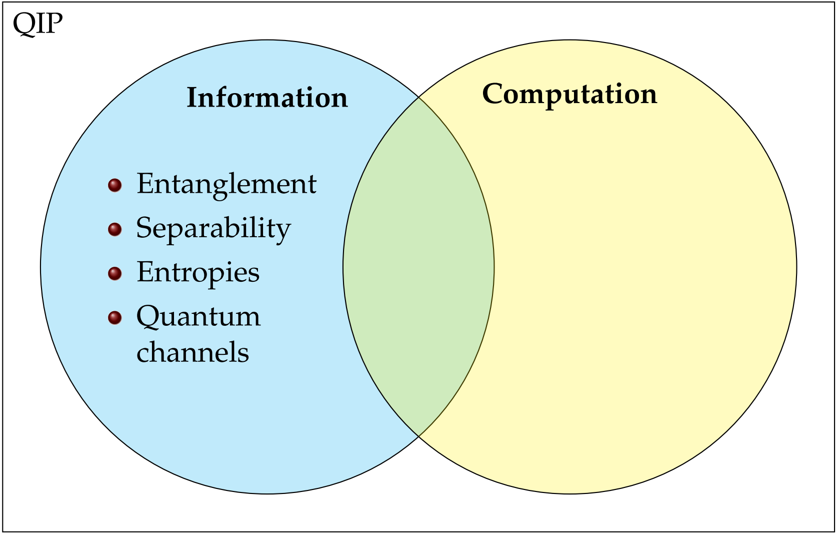
QIP
Information
Computation
Entanglement
Separability
Entropies
Quantum channels

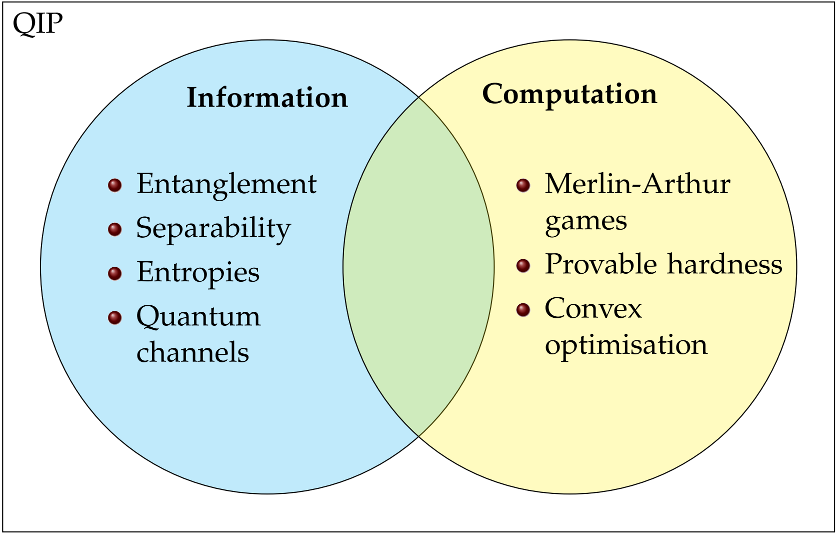
QIP
Information
Entanglement
Separability
Entropies
Quantum channels
Computation
Merlin-Arthur games
Provable hardness
Convex optimisation

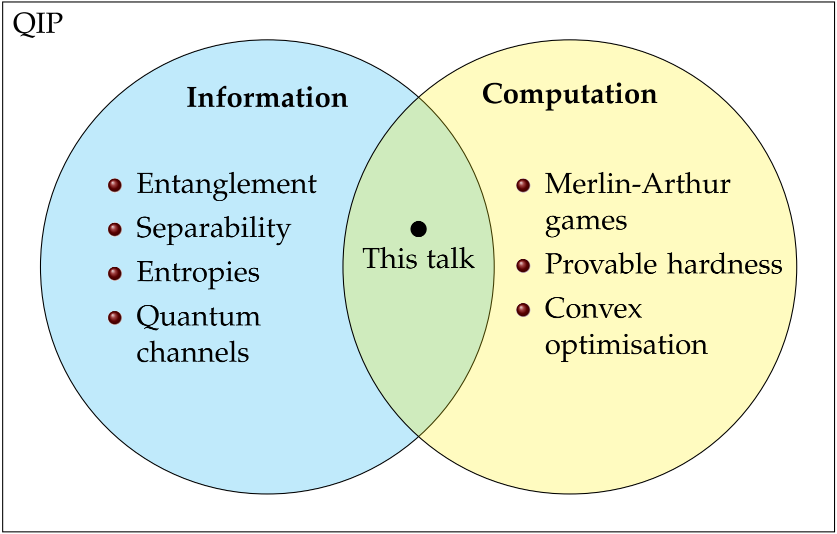
QIP
Information
Entanglement
Separability
Entropies
Quantum channels
This talk
Computation
Merlin-Arthur games
Provable hardness
Convex optimisation

The basic problem

Given a quantum state, is it entangled?

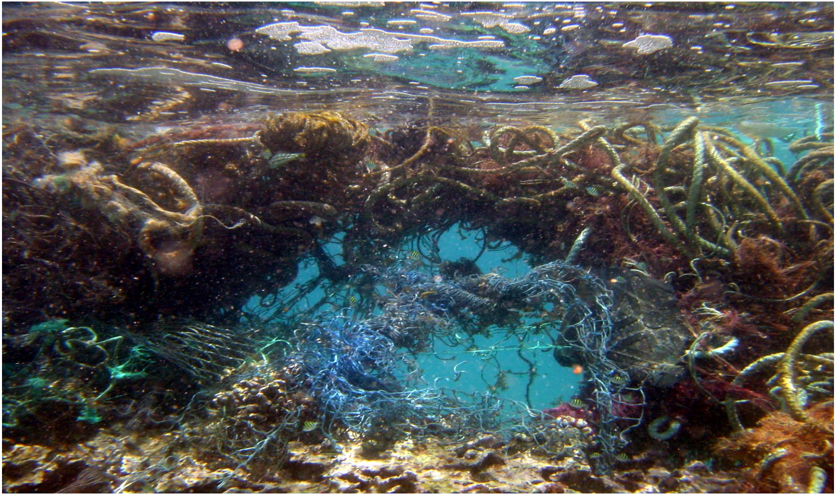

Variants

How are we given the input state?

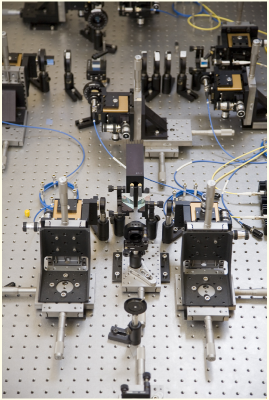

vs. $\rho = \begin{pmatrix} 1/2 & 0 & 0 & 1/2 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 1/2 \end{pmatrix}$

Variants

Is the input state pure or mixed?

$|\psi\rangle \stackrel{?}{=} |\psi_1\rangle \dots |\psi_k\rangle$ vs. $\rho \stackrel{?}{=} \sum_i p_i |\psi_1^i\rangle\langle\psi_1^i| \otimes \cdots \otimes |\psi_k^i\rangle\langle\psi_k^i|$

Is the input state bipartite or multipartite?

1 2 vs. 1 2 3 ··· k

Variants

What level of accuracy do we demand?

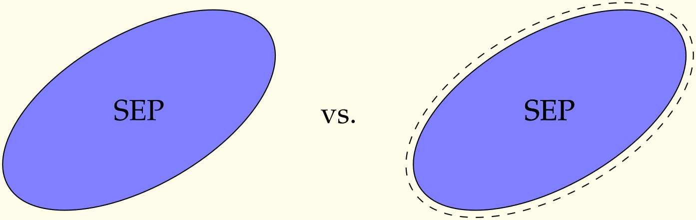

Separability testing up to accuracy ϵ: given ρ such that either $\rho \in \mathrm{SEP}$ or $\min_{\sigma \in \mathrm{SEP}} \|\rho - \sigma\|_p \geqslant \epsilon$, decide which is the case.

Variants

Do we want to detect entanglement in all states, or just some of them?

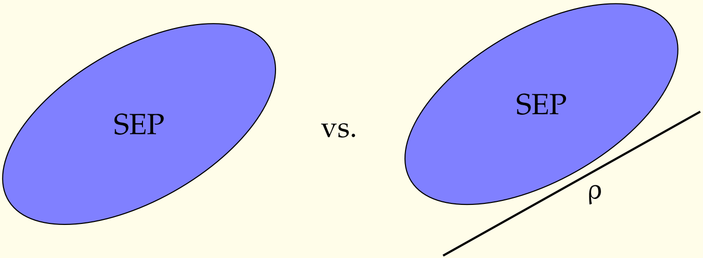

Good news and bad news

- Given a bipartite pure state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ as a d^2-dimensional vector, whether $|\psi\rangle$ is entangled can be determined efficiently using the Schmidt decomposition.

Good news and bad news

- Given a bipartite pure state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ as a d^2-dimensional vector, whether $|\psi\rangle$ is entangled can be determined efficiently using the Schmidt decomposition.

- Given a bipartite mixed state $\rho \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ as a d^2-dimensional square matrix, it's NP-hard to determine whether ρ is separable (up to accuracy $1/\mathrm{poly}(d)$).

Good news and bad news

- Given a bipartite pure state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ as a d^2-dimensional vector, whether $|\psi\rangle$ is entangled can be determined efficiently using the Schmidt decomposition.
- Given a bipartite mixed state $\rho \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ as a d^2-dimensional square matrix, it's NP-hard to determine whether ρ is separable (up to accuracy $1/\mathrm{poly}(d)$).
    - This was shown by [Gurvits '03] for accuracy $1/\exp(d)$ via a reduction from the NP-hard Clique problem.
    - Later improved to $1/\mathrm{poly}(d)$ by [Gharibian '10] (using techniques of [Liu '07]) and also (implicitly) by [Beigi '08].

Good news and bad news

- Given a bipartite pure state $|\psi\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ as a d^2-dimensional vector, whether $|\psi\rangle$ is entangled can be determined efficiently using the Schmidt decomposition.

- Given a bipartite mixed state $\rho \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ as a d^2-dimensional square matrix, it's NP-hard to determine whether ρ is separable (up to accuracy $1/\text{poly}(d)$).
  - This was shown by [Gurvits '03] for accuracy $1/\exp(d)$ via a reduction from the NP-hard Clique problem.
  - Later improved to $1/\text{poly}(d)$ by [Gharibian '10] (using techniques of [Liu '07]) and also (implicitly) by [Beigi '08].

- Stop press: There's an $\exp(O(\epsilon^{-2} \log^2 d))$ algorithm for testing separability up to accuracy ϵ in the 2-norm [Brandão, Christandl, Yard '10]!

Main result

Theorem

Let $|\psi\rangle \in (\mathbb{C}^d)^{\otimes k}$ be a pure k-partite state such that the nearest product state to $|\psi\rangle$ is the state $|\phi_1\rangle \dots |\phi_k\rangle$, where $|\langle\psi|\phi_1, \dots, \phi_k\rangle|^2 = 1 - \epsilon$.

Then there is an efficient quantum test, called the **product test**, that accepts with probability $1 - \Theta(\epsilon)$, given two copies of $|\psi\rangle$.

Main result

Theorem

Let $|\psi\rangle \in (\mathbb{C}^d)^{\otimes k}$ be a pure k-partite state such that the nearest product state to $|\psi\rangle$ is the state $|\phi_1\rangle \ldots |\phi_k\rangle$, where $|\langle\psi|\phi_1, \ldots, \phi_k\rangle|^2 = 1 - \epsilon$.

Then there is an efficient quantum test, called the **product test**, that accepts with probability $1 - \Theta(\epsilon)$, given two copies of $|\psi\rangle$.

Some notes:

- The bounds on acceptance probability don't depend on the local dimension d or the number of subsystems k.
- This is similar to classical property testing algorithms.
- The test can also be used to determine if a unitary operator is a tensor product.

The swap test

The product test uses the swap test as a subroutine.

Swap test [Buhrman et al '01]

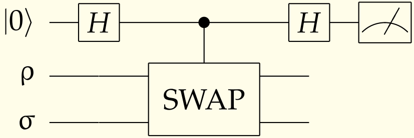

The swap test

The product test uses the swap test as a subroutine.

Swap test [Buhrman et al '01]

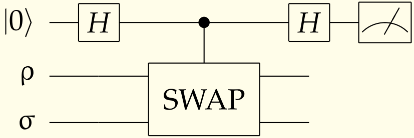

This test takes two (possibly mixed) states ρ, σ as input, returning 0 ("same") with probability

$$\frac{1}{2} + \frac{1}{2}\,\mathrm{tr}(\rho\,\sigma),$$

otherwise returning 1 ("different").

The product test

Product test

1. Prepare two copies of $|\psi\rangle \in (\mathbb{C}^d)^{\otimes k}$; call these $|\psi_1\rangle, |\psi_2\rangle$.
2. Perform the swap test on each of the k pairs of corresponding subsystems of $|\psi_1\rangle, |\psi_2\rangle$.
3. If all of the tests returned "same", accept. Otherwise, reject.

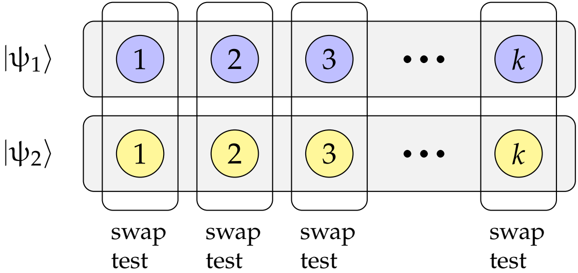

Previous use of the product test

The product test has appeared before in the literature.

- Originally introduced by [Mintert, Kuś, Buchleitner '05] as one of a family of tests for generalisations of the concurrence entanglement measure.

Previous use of the product test

The product test has appeared before in the literature.

- Originally introduced by [Mintert, Kuś, Buchleitner '05] as one of a family of tests for generalisations of the concurrence entanglement measure.
- Implemented experimentally for bipartite states by [Walborn et al '06].

Previous use of the product test

The product test has appeared before in the literature.

- Originally introduced by [Mintert, Kuś, Buchleitner '05] as one of a family of tests for generalisations of the concurrence entanglement measure.
- Implemented experimentally for bipartite states by [Walborn et al '06].
- Proposed by [AM, Osborne '08] as a means of determining whether a unitary operator is product.

Previous use of the product test

The product test has appeared before in the literature.

- Originally introduced by [Mintert, Kuś, Buchleitner '05] as one of a family of tests for generalisations of the concurrence entanglement measure.
- Implemented experimentally for bipartite states by [Walborn et al '06].
- Proposed by [AM, Osborne '08] as a means of determining whether a unitary operator is product.

Our contribution: to prove correctness of the test for all k.

Analysing the product test

Lemma

Let $P_{\text{test}}(\rho)$ be the probability that the product test passes on input ρ. Then

$$P_{\text{test}}(\rho) = \frac{1}{2^k} \sum_{S \subseteq [k]} \operatorname{tr} \rho_S^2.$$

Analysing the product test

Lemma

Let $P_{\text{test}}(\rho)$ be the probability that the product test passes on input ρ. Then

$$P_{\text{test}}(\rho) = \frac{1}{2^k} \sum_{S \subseteq [k]} \operatorname{tr} \rho_S^2.$$

- Thus the product test measures the average purity of ρ across bipartitions.
- It's immediate that $P_{\text{test}}(\rho) = 1$ if and only if ρ is a pure product state.
- Our main result says: if the average entanglement across bipartitions of $|\psi\rangle$ is low, $|\psi\rangle$ must in fact be close to a product state.

Details of main result

Theorem

Let the nearest product state to $|\psi\rangle$ be $|\phi_1\rangle \dots |\phi_k\rangle$, and set $|\langle\psi|\phi_1, \dots, \phi_k\rangle|^2 = 1 - \epsilon$. Then

$$1 - 2\epsilon + \epsilon^2 \leqslant P_{\text{test}}(|\psi\rangle\langle\psi|) \leqslant 1 - \epsilon + \epsilon^{3/2} + \epsilon^2.$$

Furthermore, if $\epsilon \geqslant 11/32$, $P_{\text{test}}(|\psi\rangle\langle\psi|) \leqslant 501/512$.

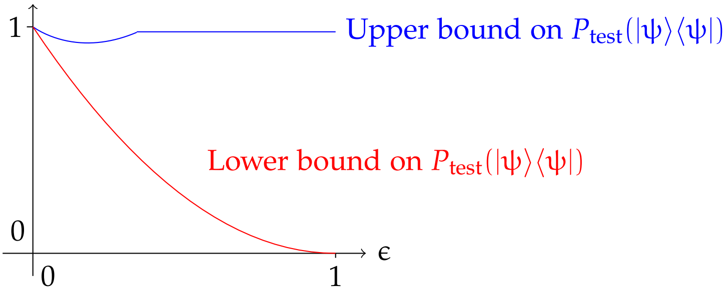

Optimality of the product test

Can we do better than the product test?

Optimality of the product test

Can we do better than the product test?

Theorem

- No "non-trivial" test can use only one copy of $|\psi\rangle$.
- The product test is "optimal" among all tests that use two copies of $|\psi\rangle$ and accept product states with certainty.

Optimality of the product test

Can we do better than the product test?

Theorem

- No "non-trivial" test can use only one copy of $|\psi\rangle$.
- The product test is "optimal" among all tests that use two copies of $|\psi\rangle$ and accept product states with certainty.

How bad is our analysis of the product test?

Optimality of the product test

Can we do better than the product test?

Theorem

- No "non-trivial" test can use only one copy of $|\psi\rangle$.
- The product test is "optimal" among all tests that use two copies of $|\psi\rangle$ and accept product states with certainty.

How bad is our analysis of the product test?

Theorem

- The leading order constants cannot be improved.
- There is a state $|\psi\rangle$ which is arbitrarily far from product and has $P_{\text{test}}(|\psi\rangle\langle\psi|) \approx 1/2$.

So (informally) these results can't be improved too much without adding dependence on k or d.

Aside: the depolarising channel

Consider the depolarising channel with noise rate $1 - \delta$, i.e.

$$\mathcal{D}_\delta(\rho) = (1-\delta)(\operatorname{tr}\rho)\frac{I}{d} + \delta\,\rho.$$

This channel's maximum output purity is multiplicative, i.e.

$$P_{\max}(\delta) := \max_\rho \operatorname{tr}(\mathcal{D}_\delta^{\otimes k}(\rho))^2$$

is achieved by product state inputs [Amosov, Holevo, Werner '00].

Aside: the depolarising channel

Consider the depolarising channel with noise rate $1 - \delta$, i.e.

$$\mathcal{D}_\delta(\rho) = (1-\delta)(\operatorname{tr}\rho)\frac{I}{d} + \delta\,\rho.$$

This channel's maximum output purity is multiplicative, i.e.

$$P_{\max}(\delta) := \max_\rho \operatorname{tr}(\mathcal{D}_\delta^{\otimes k}(\rho))^2$$

is achieved by product state inputs [Amosov, Holevo, Werner '00]. It turns out that

$$\operatorname{tr}(\mathcal{D}_\delta^{\otimes k}(\rho))^2 \propto \sum_{S \subseteq [k]} \gamma^{|S|} \operatorname{tr}\rho_S^2,$$

for some constant γ depending on δ and d.

Aside: the depolarising channel

Consider the depolarising channel with noise rate $1-\delta$, i.e.

$$\mathcal{D}_\delta(\rho) = (1-\delta)(\operatorname{tr}\rho)\frac{I}{d} + \delta\,\rho.$$

This channel's maximum output purity is multiplicative, i.e.

$$P_{\max}(\delta) := \max_\rho \operatorname{tr}(\mathcal{D}_\delta^{\otimes k}(\rho))^2$$

is achieved by product state inputs [Amosov, Holevo, Werner '00]. It turns out that

$$\operatorname{tr}(\mathcal{D}_\delta^{\otimes k}(\rho))^2 \propto \sum_{S\subseteq[k]} \gamma^{|S|} \operatorname{tr}\rho_S^2,$$

for some constant γ depending on δ and d. We have

Theorem

For small enough δ, if $\operatorname{tr}(\mathcal{D}_\delta^{\otimes k}|\psi\rangle\langle\psi|)^2 \geqslant (1-\epsilon)P_{\max}(\delta)$, there is a product state $|\phi_1,\dots,\phi_k\rangle$ with $|\langle\psi|\phi_1,\dots,\phi_k\rangle|^2 \geqslant 1-O(\epsilon)$.

This is a stability result for this channel.

Quantum Merlin-Arthur games

The complexity class QMA is the quantum analogue of NP.

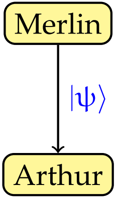

- Arthur has some decision problem of size n to solve, and Merlin wants to convince him that the answer is "yes".
- Merlin sends him a quantum state $|\psi\rangle$ of poly(n) qubits. Arthur runs some polynomial-time quantum algorithm $\mathcal{A}$ on $|\psi\rangle$ and his input and outputs "yes" if the algorithm says "accept".

Quantum Merlin-Arthur games

We say that the language L (where L is the set of bit strings we want to accept) is in QMA if there is an $\mathcal{A}$ such that, for all x:

- **Completeness:** If $x \in L$, there exists a witness $|\psi\rangle$, a state of $\text{poly}(n)$ qubits, such that $\mathcal{A}$ outputs "accept" with probability at least $2/3$ on input $|x\rangle |\psi\rangle$.

- **Soundness:** If $x \notin L$, then $\mathcal{A}$ outputs "accept" with probability at most $1/3$ on input $|x\rangle |\psi\rangle$, for all states $|\psi\rangle$.

The constants $1/3$ and $2/3$ can be amplified to be exponentially close to 0 and 1, respectively, using (e.g.) parallel repetition.

Quantum Merlin-Arthur games

QMA(k) is a variant where Arthur has access to k unentangled Merlins.

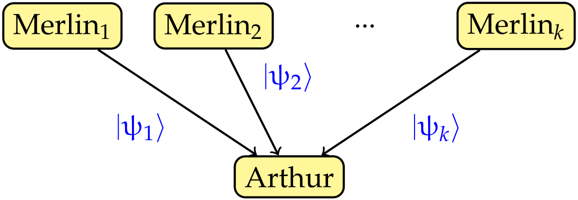

This might be more powerful than QMA because the lack of entanglement helps Arthur tell when the Merlins are cheating.

Quantum Merlin-Arthur games

A language L is in $\mathsf{QMA}(k)_{s,c}$ if there's an $\mathcal{A}$ such that, for all x:

- **Completeness:** If $x \in L$, there exist k witnesses $|\psi_1\rangle, \dots, |\psi_k\rangle$, each a state of $\mathrm{poly}(n)$ qubits, such that $\mathcal{A}$ outputs "accept" with probability at least c on input $|x\rangle |\psi_1\rangle \dots |\psi_k\rangle$.

- **Soundness:** If $x \notin L$, then $\mathcal{A}$ outputs "accept" with probability at most s on input $|x\rangle |\psi_1\rangle \dots |\psi_k\rangle$, for all states $|\psi_1\rangle, \dots, |\psi_k\rangle$.

Also define $\mathsf{QMA}_m(k)_{s,c}$ to indicate that $|\psi_1\rangle, \dots, |\psi_k\rangle$ each involve m qubits, and write $\mathsf{QMA}(k)$ to denote $s = 1/3$, $c = 2/3$.

We need this definition because straightforward parallel repetition of $\mathsf{QMA}(k)$ protocols does not work!

$\mathrm{QMA}(k)$ as an optimisation problem

Closely related to $\mathrm{QMA}_m(k)_{s,c}$

Given a 2^{km}-dimensional matrix M with $0 \leqslant M \leqslant I$, determine whether

$$\max_{|\psi\rangle = |\psi_1\rangle \otimes \cdots \otimes |\psi_k\rangle} \langle\psi| M |\psi\rangle$$

is $\geqslant c$ or $\leqslant s$.

QMA(k) as an optimisation problem

Closely related to QMA$_m(k)_{s,c}$

Given a 2^{km}-dimensional matrix M with $0 \leqslant M \leqslant I$, determine whether

$$\max_{|\psi\rangle = |\psi_1\rangle \otimes \cdots \otimes |\psi_k\rangle} \langle\psi| M |\psi\rangle$$

is $\geqslant c$ or $\leqslant s$.

- For $k = 1$, this is an eigenvalue problem with an $\exp(m)$-time algorithm.
- For $k = 2$, we need to compute

$$\max_{\rho \in \mathrm{SEP}} \operatorname{tr} M\rho.$$

- No $\exp(m)$ time algorithm is known, and even QMA$_{\log}(2)$ is not known to be in BQP.
- Compare QMA$_{\log}$ = BQP [Marriott, Watrous '05].

A potted history of QMA(k)

2003 Kobayashi, Matsumoto and Yamakami define QMA(k).

2006 Liu, Christandl and Verstraete give a problem in QMA(2) not known to be in QMA.

2007 Blier and Tapp show that graph 3-colourability can be verified by a QMA(2) protocol with messages of length $O(\log n)$, perfect completeness, and soundness $1 - 1/\operatorname{poly}(n)$.

2008 Aaronson et al show that 3-SAT on n clauses is in $\mathrm{QMA}_{O(\log n)}(\sqrt{n}\,\mathrm{polylog}(n))_{\Omega(1),1}$.

2008 Beigi improves gap in Blier-Tapp result to $\Omega(1/n^3)$.

Replacing k Merlins with 2 Merlins

$\boxed{\text{Proof}_1}$ $\boxed{\text{Proof}_2}$ $\boxed{\text{Proof}_3}$ ... $\boxed{\text{Proof}_k}$

We would like to combine these k proofs into one proof.

Replacing k Merlins with 2 Merlins

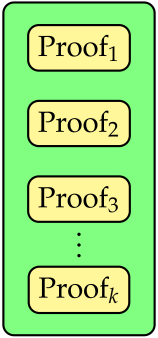

Problem: Merlin can cheat by using entanglement across proofs.

Replacing k Merlins with 2 Merlins

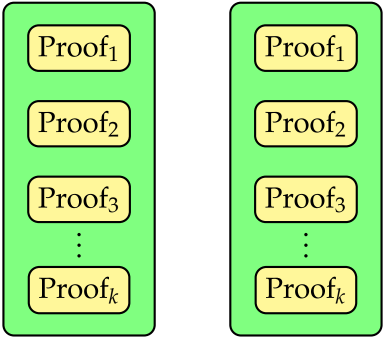

Idea: Given two copies of the proofs, we can ensure they are product states using the product test!

Then we just run the original verification algorithm on one copy.

Replacing k Merlins with 2 Merlins

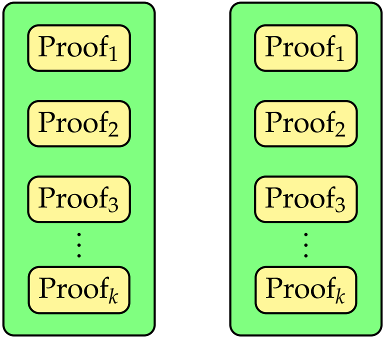

This implies that k Merlins can be simulated by 2 Merlins, up to constant soundness.

Amplification of QMA(k) protocols

- In fact, our protocol gives us something more: it turns out that the "accept" measurement operator of the new QMA(2) protocol we have produced is separable!

Amplification of $\mathrm{QMA}(k)$ protocols

- In fact, our protocol gives us something more: it turns out that the "accept" measurement operator of the new $\mathrm{QMA}(2)$ protocol we have produced is separable!
- This means that getting "accept" outcomes cannot induce entanglement between residual proofs.

Amplification of QMA(k) protocols

- In fact, our protocol gives us something more: it turns out that the "accept" measurement operator of the new QMA(2) protocol we have produced is separable!
- This means that getting "accept" outcomes cannot induce entanglement between residual proofs.
- And this means that we can amplify the soundness error to become exponentially small simply by parallel repetition of the QMA(2) protocol.

Amplification of $\mathrm{QMA}(k)$ protocols

- In fact, our protocol gives us something more: it turns out that the "accept" measurement operator of the new $\mathrm{QMA}(2)$ protocol we have produced is separable!

- This means that getting "accept" outcomes cannot induce entanglement between residual proofs.

- And this means that we can amplify the soundness error to become exponentially small simply by parallel repetition of the $\mathrm{QMA}(2)$ protocol.

- Thus, for any $k \geqslant 2$, and any c, s such that $c - s \geqslant 1/\mathrm{poly}(n)$, $\mathrm{QMA}(k)_{s,c} = \mathrm{QMA}(2)_{\exp(-n),1-\exp(-n)}$.

Amplification of $\mathrm{QMA}(k)$ protocols

- In fact, our protocol gives us something more: it turns out that the "accept" measurement operator of the new $\mathrm{QMA}(2)$ protocol we have produced is separable!

- This means that getting "accept" outcomes cannot induce entanglement between residual proofs.

- And this means that we can amplify the soundness error to become exponentially small simply by parallel repetition of the $\mathrm{QMA}(2)$ protocol.

- Thus, for any $k \geqslant 2$, and any c, s such that $c - s \geqslant 1/\operatorname{poly}(n)$, $\mathrm{QMA}(k)_{s,c} = \mathrm{QMA}(2)_{\exp(-n),1-\exp(-n)}$.

- In particular, for any $k \geqslant 2$, $\mathrm{QMA}(k) = \mathrm{QMA}(2)$.

From $\mathrm{QMA}(2)$ to hardness results

Theorem **[Aaronson et al '08]**

$$\text{3-SAT} \in \mathrm{QMA}_{\log}(\sqrt{n}\,\mathrm{polylog}(n))_{\Omega(1),1}.$$

- Our results show that satisfiability of 3-SAT formulae with n clauses can be verified by a quantum algorithm with constant success probability, given two unentangled proofs of length $O(\sqrt{n}\,\mathrm{polylog}(n))$ qubits each.

From QMA(2) to hardness results

Theorem [Aaronson et al '08]

$$3\text{-SAT} \in \mathrm{QMA}_{\log}(\sqrt{n}\,\mathrm{polylog}(n))_{\Omega(1),1}.$$

- Our results show that satisfiability of 3-SAT formulae with n clauses can be verified by a quantum algorithm with constant success probability, given two unentangled proofs of length $O(\sqrt{n}\,\mathrm{polylog}(n))$ qubits each.
- So imagine we could estimate the success probability of a QMA(2) protocol that uses proofs of dimension d, up to a constant, in time $\mathrm{poly}(d)$.
- Then this would give a subexponential-time ($2^{O(\sqrt{n}\,\mathrm{polylog}(n))}$) algorithm for 3-SAT!

From QMA(2) to hardness results

Theorem [Aaronson et al '08]

$$\text{3-SAT} \in \text{QMA}_{\log}(\sqrt{n}\,\text{polylog}(n))_{\Omega(1),1}.$$

- Our results show that satisfiability of 3-SAT formulae with n clauses can be verified by a quantum algorithm with constant success probability, given two unentangled proofs of length $O(\sqrt{n}\,\text{polylog}(n))$ qubits each.
- So imagine we could estimate the success probability of a QMA(2) protocol that uses proofs of dimension d, up to a constant, in time $\text{poly}(d)$.
- Then this would give a subexponential-time ($2^{O(\sqrt{n}\,\text{polylog}(n))}$) algorithm for 3-SAT!

So we can show hardness results for QMA(2), based on the assumption that this isn't possible.

Example hardness results

Problem OPT

Given a matrix $M \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ with $0 \leqslant M \leqslant I$, estimate

$$\max_{\rho \in \mathrm{SEP}} \operatorname{tr} M\rho$$

up to additive error δ.

Example hardness results

Problem OPT

Given a matrix $M \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ with $0 \leqslant M \leqslant I$, estimate

$$\max_{\rho \in \mathrm{SEP}} \operatorname{tr} M\rho$$

up to additive error δ.

There is a constant $\delta > 0$ such that, if 3-SAT on n clauses can't be solved in:

- ...time $\exp(\sqrt{n}\,\mathrm{polylog}(n))$, there is no $\mathrm{poly}(d)$-time algorithm for OPT.
- ...time $\exp(o(n))$, there is no $d^{O(\log^{1-\epsilon} d)}$-time algorithm for OPT, for any $\epsilon > 0$.

Problems at least as hard as OPT

Estimating minimum output entropies of quantum channels

For a quantum channel $\mathcal{N}$, determine

$$S^{\min}(\mathcal{N}) := \min_{\rho} S(\mathcal{N}(\rho))$$

up to a constant, where $S(\rho) = -\operatorname{tr} \rho \log \rho$. Also holds for estimating all Rényi entropies.

Estimating capacities of quantum channels [Beigi, Shor '08]

Estimate the Holevo capacity of $\mathcal{N}$, defined as

$$\chi(\mathcal{N}) := \max_{p_i, \rho_i} S\Big(\sum_i p_i \mathcal{N}(\rho_i)\Big) - \sum_i p_i S(\mathcal{N}(\rho_i)).$$

Problems at least as hard as OPT

Estimating ground state energies of mean-field Hamiltonians [Fannes, Vandenplas '06]

For $M \in \mathcal{B}(\mathbb{C}^d \otimes \mathbb{C}^d)$ with $0 \leqslant M \leqslant I$, define $H \in \mathcal{B}((\mathbb{C}^d)^{\otimes n})$ by

$$H = \frac{1}{n(n-1)} \sum_{1 \leqslant i \neq j \leqslant n} I - M^{(i,j)}.$$

Estimate the ground state energy of H ($\approx 1 - \max_{\rho \in \mathrm{SEP}} \operatorname{tr} M\rho$).

Determining membership in convex sets that approximate the set of separable states

Let S be a convex set approximating SEP up to Hausdorff distance δ, i.e.

$$\max\{\sup_{\rho \in S} \inf_{\sigma \in \mathrm{SEP}} \|\rho - \sigma\|_1, \sup_{\rho \in \mathrm{SEP}} \inf_{\sigma \in S} \|\rho - \sigma\|_1\} \leqslant \delta.$$

Determine membership in S.

Conclusions

- The product test is an efficient test for pure product states of n quantum systems.
- Testing pure-state entanglement is easy, so testing mixed-state entanglement is hard.
- 2 Merlins are "as good as" k Merlins: $\mathrm{QMA}(k) = \mathrm{QMA}(2)$ for $k \geqslant 2$.
- Quantum information theory and quantum computation are intimately linked.

Open problems

- Improve the best known bounds on QMA(2). Currently all we know is $\mathrm{QMA} \subseteq \mathrm{QMA}(2) \subseteq \mathrm{NEXP}$!
- What is the power of $\mathrm{QMA}(k)$ where the verifier is restricted to LOCC measurements?
  - Brandão, Christandl and Yard: $\mathrm{QMA}_{\mathrm{LOCC}}(k) = \mathrm{QMA}$ for constant k, but...
  - ...Chen and Drucker: $\mathrm{QMA}_{\mathrm{LOCC}}(\tilde{O}(\sqrt{n}))$ has efficient proofs for 3-SAT.
- Remove the convexity requirement in our "hardness of separability testing" result.
- Tighten our analysis of the product test.
- Prove stability for other channels and other Rényi entropies.
- Find other quantum property testers.

Advertisement

- There is a 2-year post-doctoral research position available at the Centre for Quantum Information and Foundations at the University of Cambridge, UK.

- The position is available now (start date flexible) and is funded by the EC FP7 project QCS ("Quantum Computer Science").

- Would suit someone with research interests in the theory of quantum computation.

- For further details, contact Prof. Richard Jozsa (rj310@cam.ac.uk).

The upper bound

The map of the first part of the proof:

- Let $|0^n\rangle$ be the closest product state to $|\psi\rangle$.
- Write $|\psi\rangle = \sqrt{1-\epsilon}\,|0^n\rangle + \sqrt{\epsilon}\,|\phi\rangle$ for some $|\phi\rangle$.
- This allows us to calculate $\sum_S \operatorname{tr}\psi_S^2$ explicitly in terms of ϵ, $|\phi\rangle$.
- Writing $|\phi\rangle = \sum_x \alpha_x |x\rangle$, can upper bound $\sum_S \operatorname{tr}\psi_S^2$ in terms of how much weight $|\phi\rangle$ has on low Hamming weight basis states.
- Showing that there can be no weight on states of Hamming weight 1 completes the proof.

The second part of the proof

The first part of the proof ends up showing

$$P_{\text{test}}(|\psi\rangle\langle\psi|) \leqslant 1 - \epsilon + \epsilon^{3/2} + \epsilon^2.$$

This bound is greater than 1 for large ϵ!

We fix up the proof by showing (roughly):

- $P_{\text{test}}(|\psi\rangle\langle\psi|)$ is upper bounded by the probability that the product test across any partition into k parties passes.
- If $|\psi\rangle$ is far from product across the n subsystems, one can find a k-partition such that the distance from the closest product state (wrt this partition) falls into the regime where the first part of the proof works.

The second part of the proof

The first part of the proof ends up showing

$$P_{\text{test}}(|\psi\rangle\langle\psi|) \leqslant 1 - \epsilon + \epsilon^{3/2} + \epsilon^2.$$

This bound is greater than 1 for large ϵ!

We fix up the proof by showing (roughly):

- $P_{\text{test}}(|\psi\rangle\langle\psi|)$ is upper bounded by the probability that the product test across any partition into k parties passes.
- If $|\psi\rangle$ is far from product across the n subsystems, one can find a k-partition such that the distance from the closest product state (wrt this partition) falls into the regime where the first part of the proof works.
- This leads to the result that, if $\epsilon \geqslant 11/32$, $P_{\text{test}}(|\psi\rangle\langle\psi|) \leqslant 501/512$.

These constants can clearly be improved somewhat...