Upper Bounds in Quantum Complexity

We've established $\mathrm{BQP} \subseteq \mathrm{PP}$. What about $\mathrm{QMA} \subseteq \mathrm{PP}$?

We could imagine looping through some discretized description of the set of all possible quantum witness states. Such a proof could work, but needs to be explicit about how to search over witness states.

Another way to show $\mathrm{QMA} \subseteq \mathrm{PP}$ would be to show that we can solve the local Hamiltonian problem in PP. This is the route we will take, using a Feynman path integral.

Given a local Hamiltonian $H = H_1 + H_2 + \ldots + H_m$ and two real numbers (a,b) with $b - a \geq 1/\mathrm{poly}(n)$ we seek to decide between two cases:

YES: $E_0(H) \leq a$

NO: $E_0(H) \geq b$

Upper Bounds in Quantum Complexity

To distinguish between these cases, define $G = I - \epsilon H$, where $\epsilon \leq \|H\|^{-1}$. The eigenvalues of G are contained in the interval [0,1], and the largest eigenvalue is $\mu = 1 - \epsilon E_0(H)$.

We will consider a "partition function" $Z = \mathrm{tr}\left(G^L\right)$.

In a YES instance we have $\mu \geq \mu^+$, and in a NO instance $\mu \leq \mu_-$, with $\mu_+ - \mu_- > 1/\mathrm{poly}(n)$

Can we upper bound Z for NO instances in terms of μ_- , and lower bound Z for YES instances with μ_+ ?

YES: $Z \geq \mu_+^L$

NO: $Z \leq 2^n \mu_-^L$

$$\Longrightarrow \qquad \frac{Z_{\mathrm{yes}}}{Z_{\mathrm{no}}} \geq 2^{\frac{L}{\mathrm{poly}(n)} - n}$$

So L = poly(n) suffices to make the ratio of partition functions between the two cases exponentially large.

Upper Bounds in Quantum Complexity

Expanding Z as a path integral:

$$Z = \mathrm{tr}\left(G^L\right) = \sum_{z_1 \in \{0,1\}^n} \langle z_1 | G^L | z_1 \rangle = \sum_{z_1,\ldots,z_L \in \{0,1\}^n} \langle z_1|G|z_2\rangle\langle z_2|G|z_3\rangle \ldots \langle z_{L-1}|G|z_L\rangle\langle z_L|G|z_1\rangle$$

More compactly: $Z = \mathrm{tr}\left(G^L\right) = \sum_{z_1,\ldots,z_L \in \{0,1\}^n} \prod_{i=1}^{L} \langle z_i|G|z_{i+1}\rangle$

We can compute the amplitudes on at a time using polynomial space. We can put all the positive amplitudes on one side of the ledger, and the negative amplitudes on the other side.

Equivalently, since the ratio $Z_{\mathrm{yes}}/Z_{\mathrm{no}}$ is exponentially large, we can decide between the two cases with probability better than random guessing, just by looking at the amplitude of a single random path.

Upper Bounds in Quantum Complexity

From the perspective of the Feynman path integral, the power of quantum computing (BQP, or even QMA) comes from the ability to approximate exponentially large alternating sums.

These path integrals add up exp many amplitudes, which may be positive or negative and thus undergo many cancellations, and the result depends on what is left over. This is **interference**.

It's a bit like the year 2000 US presidential election: +50 million votes for Bush, +50 million votes for Gore, and Bush won because he led by 537 votes in the Florida recount.

Aaronson: "as far as I'm concerned, Feynman got the Nobel Prize for showing $\mathrm{BQP} \subseteq PP$." …

The close ties between quantum complexity and alternating sums raises a question. What is the complexity of quantum systems described by path integrals with all positive amplitudes?

Upper Bounds in Quantum Complexity

Returning to the "partition function": $Z = \mathrm{tr}\left(G^L\right) = \sum_{z_1,\ldots,z_L \in \{0,1\}^n} \prod_{i=1}^{L} \langle z_i|G|z_{i+1}\rangle$

We could put aside H for a moment, and rephrase the local Hamiltonian problem as "deciding the largest eigenvalue of a Hermitian matrix."

What happens to the complexity of estimating this sum if all the terms have a nonnegative amplitudes?

$$\prod_{i=1}^{L} \langle z_i|G|z_{i+1}\rangle \geq 0 \quad \forall z_1, \ldots, z_L$$

A sufficient condition for the amplitudes to all be nonnegative would be for G to be a matrix with nonnegative entries (in the basis of the $|z_i\rangle$).

In what follows we will (1) motivate the case of nonnegative G in terms of the Hamiltonian, and (2) establish a protocol with a classical prover and classical verifier for solving this restricted version of LH.

Upper Bounds in Quantum Complexity

An example of a Hamiltonian giving rise to a nonnegative matrix G, consider the Ising model:

$$H = -\sum_{i=1}^{n} X_i - \sum_{i=1}^{n} Z_i Z_{i+1}$$

(one may specify *this* Ising model as the “ferromagnetic transverse Ising spin chain”)

We want to understand H as a matrix in the computational basis, and show that it can shifted and rescaled into a nonnegative matrix G.

The ZZ terms are on the diagonal, so we can control the sign of those (make them all positive or negative) by shifting the energy by a constant. $H \mapsto H - E_0 I$

What about the off-diagonal matrix elements of H in the computational basis?

Upper Bounds in Quantum Complexity

The off-diagonal elements of $\sum_i X_i$ are (real and) nonnegative, so the off-diagonal elements of H are real and non-positive.

Therefore, after shifting the ground energy to zero, $G = I - H/\|H\|$ is a nonnegative matrix.

Any H with all real and non-positive matrix elements in a basis B can be rescaled and shifted into a nonnegative matrix in that basis. Physicist say Hamiltonians with this property "do not have a sign problem", and in quantum information we call these Hamiltonians **stoquastic**.

"Stoquastic" = "Quantum" + "Stochastic." These Hamiltonians have some properties in common with stochastic matrices. The key point is that their ground state path integrals are nonnegative sums, instead of alternating sums, and this will limit the complexity of the stoquastic local Hamiltonian problem.

Upper Bounds in Quantum Complexity

Before using a path integral to solve the stoquastic local Hamiltonian problem using a classical prover and classical verifier, it's worth noting that stoquastic Hamiltonians are ubiquitous in nature.

Besides the ferromagnetic transverse Ising chain, all generalized transverse Ising models are stoquastic:

$$H = -\sum_{i=1}^{n} X_i - \sum_{i,j} \alpha_{ij} Z_i Z_j$$

The ferromagnetic Heisenberg model is also stoquastic. $H = -\sum_{i\ j} \vec{S}_i \cdot \vec{S}_j$

In physics, any H = T + U (kinetic + potential) for distinguishable particles (or bosons) is stoquastic e.g. particles hopping on a graph under the influence of a potential:

$$H = -\sum_{(u,v)\in E} (|u\rangle\langle v| + |v\rangle\langle u|) + \sum_{v\in V} U_v |v\rangle\langle v|$$

Upper Bounds in Quantum Complexity

Before using a path integral to solve the stoquastic local Hamiltonian problem using a classical prover and classical verifier, it's worth noting that stoquastic Hamiltonians are ubiquitous in nature.

Besides the ferromagnetic transverse Ising chain, all generalized transverse Ising models are stoquastic:

$$H = -\sum_{i=1}^{n} X_i - \sum_{i,j} \alpha_{ij} Z_i Z_j$$

The ferromagnetic Heisenberg model is also stoquastic. $H = -\sum_{i\ j} \vec{S}_i \cdot \vec{S}_j$

In physics, any H = T + U (kinetic + potential) for distinguishable particles (or bosons) is stoquastic e.g. particles hopping on a graph under the influence of a potential:

$$H = -\sum_{(u,v)\in E} (|u\rangle\langle v| + |v\rangle\langle u|) + \sum_{v\in V} U_v |v\rangle\langle v|$$

Upper Bounds in Quantum Complexity

In contrast with the general local Hamiltonian problem which is QMA-complete, the stoquastic LH problem is in the classical class AM. This is proven in "The complexity of stoquastic local Hamiltonian problems". Bravyi, Divincenzo, Oliviera, Terhal, 2006.

We have already defined the class MA: Merlin sends a poly sized classical witness to Arthur, who verifies it using a BPP machine. AM is similar, but now Arthur goes first: he can flip some coins and ask some questions before Merlin sends the witness.

The Arthur-Merlin protocol for deciding stoquastic LH is based on approximating the partition function:

$$Z = \mathrm{tr}\left(G^L\right) = \sum_{z_1,\ldots,z_L \in \{0,1\}^n} \prod_{i=1}^{L} \langle z_i | G | z_{i+1} \rangle$$

The proof illustrates a general theme in approximation algorithms called "the connection between sampling and counting."

Upper Bounds in Quantum Complexity

The first step is to reduce our sum over nonnegative weights to an unweighted counting problem.

$$Z = \mathrm{tr}\left(G^L\right) = \sum_{z_1,\ldots,z_L \in \{0,1\}^n} \prod_{i=1}^{L} \langle z_i | G | z_{i+1} \rangle$$

Since the matrix elements of G are specified by m = poly(n) many bits, we can decompose G as an average:

$$G = \frac{1}{2^m} \sum_{t \in \{0,1\}^m} G(t)$$

where each G(t) is a binary matrix (a matrix with entries 0 and 1). This decomposition is highly nonunique, there are many explicit choices one could make (it's also inessential, but used in Bravyi et al.'s proof).

Upper Bounds in Quantum Complexity

Now the trace is re-expressed as:

$$\mathrm{tr}\left(G^L\right) = \frac{1}{2^{mL}} \sum_{t_1,\ldots,t_L \in \{0,1\}^m} \mathrm{tr}\left(G(t_1)\ldots G(t_L)\right) = \frac{1}{2^{mL}} \sum_{s \in \{0,1\}^{mL}} F(s)$$

Where $s = (t_1, ..., t_L, z_1, ..., z_L)$ is a binary string of length (n + m)L, and F is a Boolean function:

$$F(s) = \langle z_1|G(t_1)|z_2\rangle\langle z_2|G(t_2)|z_3\rangle \ldots \langle z_L|G(t_L)|z_1\rangle \in \{0,1\}$$

Therefore the original problem is reduced to counting the cardinality of the support of F,

$$Z = \mathrm{tr}\left(G^L\right) = \frac{1}{2^{mL}} \sum_{s \in \{0,1\}^{mL}} F(s) = \mathrm{supp}(F) = 2^{-mL}|\Omega| \quad , \quad \Omega = \{x \in \{0,1\}^{(n+m)L} : F(x) \neq 0\}$$

Upper Bounds in Quantum Complexity

Therefore the problem of deciding whether the ground state energy is large or small is now reduced to deciding whether Ω is large of small. As Bravyi et al. phrase it:

$$\begin{aligned} \mu(G) \geq \mu_+ \quad &\Longrightarrow \quad |\Omega| \geq \text{LARGE} \\ \mu(G) \leq \mu_- \quad &\Longrightarrow \quad |\Omega| < \text{SMALL}, \end{aligned}$$

where

$$\text{LARGE} = 2^{L(p_2(n)+\log \mu_+)} \quad \text{and} \quad \text{SMALL} = 2^{L(p_2(n)+\log \mu_- + \frac{n}{L})},$$

such that

$$\text{LARGE} = 2^n \cdot \text{SMALL} \quad \text{if} \quad L = 2np_1(n).$$

Thus it suffices for Merlin to prove a lower bound $|\Omega| \geq \text{LARGE}$.

Immediately after this line, the proof is concluded with the following argument.

Upper Bounds in Quantum Complexity

We can now invoke the Goldwasser and Sipser approximate counting protocol [20] based on Carter-Wegman universal hashing functions [21]. Recall that Ω is a set of k-bit strings, where $k = L(n + p_2(n))$. The main idea of [20] is that Arthur can compress k-bit strings to shorter b-bit strings using randomly chosen linear hash functions. One can choose parameters of the hashing such that the image $h(\Omega) \subseteq \Sigma^b$ is sufficiently dense (for positive instances). Arthur estimates the volume of $h(\Omega)$ using the standard Monte-Carlo method: he generates a large list of random b-bit strings and estimates the fraction of strings that belong to $h(\Omega)$. At this stage he needs Merlin's help, since a membership in the set $h(\Omega)$ is no longer efficiently verifiable because each string in Σ^b may have exponentially large number of pre-images. On the other hand, Merlin can prove a membership in the set $h(\Omega)$ by sending Arthur any of pre-images. In Appendix A we give some details of the parameters of the hash functions.

Upper Bounds in Quantum Complexity

This proves that the stoquastic local Hamiltonian problem can be solved in AM. In contrast, the general LH problem is QMA-complete, so in this way stoquastic H have less ground state complexity.

But what about computational physics, can we use the nonnegativity of stoquastic path integrals to simulate some quantum systems efficiently, even without the help of a powerful prover?

The answer turns out to be yes: because stoquastic path integrals are nonnegative, we can define a distribution on the space of paths. We can use a Markov chain to sample from that distribution, and then use those samples to compute physical observables by the Monte Carlo method.

Methods which use this approach are called quantum Monte Carlo methods. The particular variant we are describing is called “path integral Monte Carlo.” It was the first and still the best in practice (for spin systems), and was proposed by Suzuki, “Monte Carlo simulations of quantum spin systems”, 1977.

Upper Bounds in Quantum Complexity

The path integrals we've defined in terms of the matrix G apply to ground states. We can get also include thermal states by working with the physical partition function:

$$Z = \mathrm{tr}\left(e^{-\beta H}\right) = \sum_{z_1 \ldots z_L} \prod_{i=1}^{L} \langle z_i | e^{-\beta H/L} | z_{i+1} \rangle \quad , \quad z_{L+1} := z_1$$

Each $e^{-\beta H/L}$ is a nonnegative matrix, and so we may define a probability distribution over paths:

$$\pi(z_1, \ldots, z_L) = \frac{1}{Z} \prod_{i=1}^{L} \langle z_i | e^{-\beta H/L} | z_{i+1} \rangle$$

The marginal distribution is: $\pi(z_1) = \sum_{z_2,\ldots,z_L} \pi(z_1, \ldots, z_L) = \langle z_1 | e^{-\beta H/L} | z_1 \rangle / Z$

Therefore, if we can sample from π we can sample from the thermal state!

Upper Bounds in Quantum Complexity

Our distribution over paths has a form resembling a classical Gibbs distribution: the probability density is easy to compute given $(z_1, ..., z_L)$, but the partition function Z is not easy to compute.

With more arithmetic, one can formally relate our path integral for Z to a classical partition function. This is called "the correspondence between quantum systems in D spatial dimensions, and classical systems in D + 1 spatial dimensions." This correspondence is well studied in 1980s physics because it relates the critical phenomena of the two models, critical exponents, etc.

Suzuki, "Relationship between d-Dimensional Quantal Spin Systems and (d+1)-Dimensional Ising Systems: Equivalence, Critical Exponents and Systematic Approximants of the Partition Function and Spin Correlations." 1976.

But the applicability of the correspondence is widely misunderstood due to sloppyness: it only applies to stoquastic Hamiltonians, and it only applies to equilibrium (thermal or ground state) path integrals.

Upper Bounds in Quantum Complexity

Anyway, back to this form that resembles a classical Gibbs distribution:

$$\pi(z_1, ..., z_L) = \frac{1}{Z} \prod_{i=1}^{L} \langle z_i | e^{-\beta H/L} | z_{i+1} \rangle$$

To sample from π, one constructs a stochastic matrix P which satisfies detailed balance

$$\pi_x P_{xy} = \pi_y P_{yx}$$

Where at this moment, x and y are shorthand for configurations of our state space, which is a space of paths $(z_1, .., z_L)$. In constructing P we choose which configurations to transition between, for example it is common to stipulate that x and y differ by flipping a single bit.

The Metropolis probabilities satisfy detailed balance: $P_{xy} \propto \min\left\{1, \frac{\pi_y}{\pi_x}\right\}$

Upper Bounds in Quantum Complexity

Therefore if H is stoquastic, its thermal path integrals are nonnegative and we can in principle sample from them by a Markov chain and use this to calculate observables by Monte Carlo.

But how long does the Markov chain take to converge? To see that this is a major issue, note that our NP-complete 3-SAT Hamiltonians are also stoquastic.

To get a general result about this rate of convergence, we should choose a premise that excludes NP-hard problems. The case of ferromagnetic systems (all spins want to align, on any graph and in any dimension), was solved by Bravyi and Gosset in 2017.

The case of simulating 1D stoquastic systems, was solved by me during my PhD. ☺

$$\mathcal{H} = -\sum_{j=1}^{n} \Gamma_j \sigma_j^x + \sum_{j=1}^{n} K_j^z \sigma_j^z + \sum_{j,k=1}^{n} K_{jk}^{zz} \sigma_j^z \sigma_k^z \quad , \quad \Gamma := \min_{j=1,\ldots,n} \Gamma_j > 0,$$

$$K_j^z \in [-1,1] \quad , \quad \text{and } |K_{jk}^{zz}| \leq |i-j|^{-(2+\xi)} \text{ for } \xi > 0.$$