Frontiers of Quantum complexity *Instructor: Anand Natarajan, Anurag Anshu*

Lecture 21: State Synthesis

Nov. 18, 2025 *Scribe: Jasper Jain*

1 Setup

So far, we've discusses complexity classes of problems with classical inputs and outputs. One could imagine other types of problems:

	Classical Input	Quantum Input
Classical Ouptut	"Given x, output y"	"Given $\|\psi\rangle^{\otimes l}$, find the parent Hamiltonian H of $\|\psi\rangle$."
Quantum Output	"Given x, prepare $\|\psi_x\rangle$."	"Given $(x, \|\psi\rangle)$, prepare $U_x \|\psi\rangle$."

Here we mean x and y to be classical data, and in the quantum input/classical output case given an example of a "quantum learning theory" problem.

Today we're going to discuss the classical input/quantum output case. The hope would be that while classical complexity theory problems are incredibly difficult, these problems might have some separate complexity theory which is more tractable.

For the "state synthesis" problems we'll consider today, this hope turns out to be false. The positive viewpoint on this means that we have some good algorithms for state synthesis!

Question: For every family of states $\{|\psi_x\rangle\}_x$, does there exist a classical oracle O such that a BQP^O machine can prepare $|\psi_x\rangle$ given x?

We formulate the question like this because it might be the case that some state $|\psi_x\rangle$ is very difficult to prepare, but all of that difficulty is wrapped up in some hard-to-compute classical function. So this roughly tells if the difficulty is in somehow quantum.

If the answer is no, then there might exist some states $|\psi_x\rangle$ that are hard to prepare even in a world where P = PSPACE. If the answer is yes, then showing a lower bound for $|\psi_x\rangle$ gives a lower bound for some classical function f.

2 "A Sequence of More and More Refined Answers"

2.1 The Bernstein-Vazirani Algorithm

Suppose we have some oracle performing $O|x\rangle = (-1)^{x\cdot y}|x\rangle$ for some hidden $y \in \{0,1\}^n$. Here x is meant to be an n-qubit computational basis state.

It turns out that it is possible to prepare $|y\rangle$ with one query. Start with the uniform superposition state $\frac{1}{\sqrt{2^n}}\sum_x |x\rangle$ and apply O, and then Hadamards to each qubit:

$$H^{\otimes n}\left(O\left(\frac{1}{\sqrt{2}}\sum_x |x\rangle\right)\right) = H^{\otimes n}\left(\frac{1}{\sqrt{2^n}}\sum_x (-1)^{x\cdot y}|x\rangle\right) = y.$$

Relating this to the state preparation task, this means that if y encodes some classical description of the $|\psi\rangle$, we can read it out with one query of the oracle. Preparing the state given this information might still be complicated, but does not require additional queries of the oracle. But our vector y will need to be exponentially long.

2.2 A Better Way, due to Grover and Rudoplh (2002?)

The result here is that for any collection $\{\psi_x\}_x$ of n-qubit states, there exists an oracle O acting on states of polynomial length such that in $2n+2$ queries one can prepare $|\psi_x\rangle$ to $1/\exp(n)$ precision.

They considered states of the form

$$\sum_{x\in\{0,1\}^n} \sqrt{p(x)}\,|x\rangle = \sum_{x\in\{0,1\}^n} \sqrt{p(x_1)p(x_2|x_1)p(x_3|x_1,x_2)\cdots}\,|x\rangle\,.$$

Suppose one had a subrouting sending

$$|x_1\cdots x_k\rangle\otimes|0\rangle\otimes|0\rangle^{\otimes m} \mapsto \sqrt{p(x_{k+1}=0|x_1,\cdots x_k)}\,|x_1,\cdots x_k,0\rangle\otimes|0\rangle^{\otimes m}+\sqrt{p(x_{k+1}=1|x_1,\cdots x_k)}\,|x_1,\cdots x_k,1\rangle\otimes|0\rangle^{\otimes m}\,,$$

where the last constituent of the tensor product is an ancilla. Then we could apply this iteratively and be done. Grover and Rudolph proceed to implement this.

Suppose we have some oracle O which given $x_1,\cdots,x_k$ ouputs

$$\theta(x_1,\cdots,x_k) := \arccos\Big(\sqrt{p(0)|x_1,\cdots,x_k}\Big).$$

Then we could send

$$\begin{aligned}|x_1,\cdots,x_k\rangle\otimes|0\rangle\otimes|0\rangle &\xrightarrow{O} |x_1,\cdots,x_k\rangle\otimes|0\rangle\otimes|\theta\rangle\\ &\xrightarrow{U} |x_1,\cdots,x_k\rangle\otimes(\cos\theta\,|0\rangle+\sin\theta\,|1\rangle)\otimes|\theta\rangle\\ &\xrightarrow{O^\dagger} |x_1,\cdots,x_k\rangle\otimes(\cos\theta\,|0\rangle+\sin\theta\,|1\rangle)\otimes|0\rangle\end{aligned}$$

Here U is some unitary which does not need involve querying O, and we use $O^\dagger$ to uncompute the ancilla. Since we do this for ecah qubit, this amounts to $2n$ queries. If we had arbitrary phases $e^{i\phi_x}$ on each term and a corresponding oracle, we could do the same step and restore the phases with two calls to the oracle. Thus the algorithm takes $2n+2$ queries in total.

The original manuscript can be found here.

2.3 Simplifying the Quantum Circuit

The simpler our circuit in BQP^O, the stronger of a lower bound we'll obtain for computing a classical function f. This motivated an article from Irani, Natarajan, Nirkhe, Rao, and Yuen concerning "search to decision reductions."

The idea here is that search problems can be reduced to decision problems simply by encoding the possible answers as sequences of bits, and repeatedly asking the decision problem "what is the next bit?" For the example of NP, this means that $\mathrm{FNP}\subset P^{NP}$, where containtment follows from

asking "is there a solution to this formula where the first bit is fixed to 0?" and so on. Similarly, we might have SearchQMA $\subset \mathrm{BQP}^{\mathrm{QMA}}$, and something like $\mathrm{P} \subset \mathrm{BQP}^{QCMA}$. To think about the first of these possible containments, we note that certainly SearchQMA $\subset \mathrm{BQP}^{Q}$ for some class Q; we want to find the smallest class such that this is true.

Natarajan et al. showed that an oracle of Aaronson and Kuperberg separates SearchQCMA and $\mathrm{BQP}^{\mathrm{QCMA}}$!

An important question is: does there exist, as with the Bernstein-Vazirani Algorithm, a 1-query algorithm for finding ground states?

One class of states which is easy to prepare with a single query of a classical oracle is the "binary phase states"

$$\frac{1}{\sqrt{2^n}} \sum_x (-1)^{f(x)} |x\rangle ,$$

where f is the oracle. It turns out that (in high dimensions?) randoms states $|\psi\rangle$ drawn from the Haar measure look like

$$\frac{1}{\sqrt{2^n}} \sum_x e^{i\theta_x} |x\rangle$$

for some phases θ_x. Fixing a state $|\psi\rangle$, if we pick a unitary U randomly from the Haar measure, then $U|\psi\rangle$ looks random as well. These observations suggest the following procedure to prepare a target ground state $|\tau\rangle$:

1. Given the ground state $|\tau\rangle$ and some unitary U, define the oracle f acting on computational basis states to be $f_{U,\tau}(|\psi\rangle) = \mathrm{sgn}\left[\mathrm{Re}[\,\langle\psi|U|\tau\rangle]\right]$.
2. Prepare
$$|\psi\rangle = \frac{1}{\sqrt{2^n}} \sum_x |x\rangle ,$$
the uniform sum over computational basis states.
3. Apply the oracle to obtain the binary phase state
$$|\phi\rangle = \sum_x (-1)^{f_{U,\tau}(x)} |x\rangle .$$
4. Return $U^\dagger \phi$.

Clearly this should give some approximation for the ground state, but much is lost by discarding the phases. Still, a performance guarantee is given by:

Fact: For all states $|\psi\rangle$ on N qubits, there exists a Clifford operator C and binary phase state $|\psi^*\rangle$ such that $\langle\psi^*|C|\psi\rangle \geq .35$, where the constant is independent of N.

Natarajan et al. make some optimizations. One is to verify the result with phase estimation. Another is to use a "more random" set of matrices $\{U_i\}_i$ to sample from. Most importantly, they reduce the error using "swap test amplification." This views the outputs $U^\dagger|\phi\rangle$ as having some signal component in the direction of τ and noise in other directions. SWAP operators can be used to cancel this noise.

2.4 Rosenthal's Improvement

Theorem (Rosenthal '23): There exists a 1-query (dirty) state synthesis algorithm with $1/\exp(N)$ error which lies in $(\mathrm{QAC}^0_f)^O$.

Here "dirty" means that the ancilla is not uncomputed. This repeatedly uses the idea that for generic states $|\psi\rangle$ and $|\phi\rangle$, there are good Clifford approximations to the equation $|\phi\rangle = U|\psi\rangle$. One starts with the same oracle idea as Natarajan et al. That is, one finds a Clifford approximation $|\phi_0\rangle$ of the desired state $|\psi\rangle$. Then one defines $|\eta_1\rangle := |\psi\rangle - i|\phi_0\rangle$, uses the Clifford appriximation on $|\eta_1\rangle$ to find $|\phi_1\rangle$, defines $|\eta_2\rangle = |\psi\rangle - |\phi_0\rangle - \beta|\phi_1\rangle$,and so on. By the end, hopefully we should have

$$|\psi\rangle - \alpha \sum_k \beta^k |\phi_k\rangle \approx 0$$

for some constant α. The key step is that the oracle can know all of the η_i ahead of time! So each step can be done in parallel, only requiring one query. Thus, "a QAC^0_f-lower bound on $|\psi\rangle$ gives a QAC^0_f lower bound on f."