A Complete Characterization of Unitary Quantum Space

Bill Fefferman (QuICS, University of Maryland)

Joint with Cedric Lin (QuICS)

Based on arXiv:1604.01384

Theory Seminar, UT Austin

Our motivation: How powerful are quantum computers with a small number of qubits?

- *Our results*: Give two natural problems *characterize* the power of quantum computation with a given bound on the number of qubits
  1. **Precise Succinct Hamiltonian** problem
  2. **Well-conditioned Matrix Inversion** problem
- These characterizations have many applications
  - **QMA** proof systems and Hamiltonian complexity
  - Classical **Logspace** complexity
  - Even connections to physics (e.g., the power of preparing **PEPS** states)

Quantum space complexity

- **BQSPACE**[k(n)] is the class of promise problems $L=(L_{yes},L_{no})$ that can be decided by a bounded error quantum algorithm acting on k(n) qubits.
  - i.e., Exists uniformly generated family of quantum circuits $\{Q_x\}_{x\in\{0,1\}^*}$ each acting on $O(k(|x|))$ qubits:
    - "If answer is yes, the circuit Q_x accepts with high probability"
      $$x \in L_{yes} \Rightarrow \langle 0^k | Q_x^\dagger | 1 \rangle \langle 1 |_{out} Q_x | 0^k \rangle \geq 2/3$$
    - "If answer is no, the circuit Q_x accepts with low probability"
      $$x \in L_{no} \Rightarrow \langle 0^k | Q_x^\dagger | 1 \rangle \langle 1 |_{out} Q_x | 0^k \rangle \leq 1/3$$
- Our results show two natural complete problems for **BQSPACE**[k(n)]
  - For any k(n) so that log(n)≤k(n)≤poly(n)
  - Our reductions use classical k(n) space and poly(n) time
- *Subtlety*: This is "unitary quantum space"
  - No intermediate measurements
  - Not known if "deferring" intermediate measurements can be done space efficiently

Quantum Merlin-Arthur

- Problems whose solutions can be verified quantumly given a quantum state as witness
- **QMA**(c,s) is the class of promise problems L=(L_{yes},L_{no}) so that:

$$x \in L_{yes} \Rightarrow \exists|\psi\rangle\ \Pr[V(x,|\psi\rangle) = 1] \geq c$$

$$x \in L_{no} \Rightarrow \forall|\psi\rangle\ \Pr[V(x,|\psi\rangle) = 1] \leq s$$

- **QMA** = **QMA**(2/3,1/3) = $\cup_{c>0}$**QMA**(c,c-1/poly)
- *k-Local Hamiltonian* problem is **QMA**-complete (when k≥2)[Kitaev '00]
  - Input: $H = \sum_{i=1}^{M} H_i$, each term H_i is k-local
  - Promise either:
    - Minimum eigenvalue λ_{min}(H) > b or λ_{min}(H) < a
    - Where b-a≥1/poly(n)
  - Which is the case?
- Generalizations of **QMA**:
  1. **PreciseQMA**=$\cup_{c>0}$**QMA**(c,c-1/exp)
  2. k-**bounded QMA**$_m$(c,s)
     - Arthur's verification circuit acts on k qubits
     - Merlin sends an m qubit witness

Characterization 1:

Precise Succinct Hamiltonian problem

The *Precise Succinct* **Hamiltonian** Problem

- Definition: "*Succinct Encoding*"
  - We say a classical Turing machine M is a *Succinct Encoding* for $2^{k(n)}$ x $2^{k(n)}$ matrix A if:
    - On input $i \in \{0,1\}^{k(n)}$, M outputs non-zero elements in i-th row of A
    - Using at most poly(n) time and k(n) space
- k(n)-**Precise Succinct Hamiltonian** problem
  - Input: Succinct Encoding of $2^{k(n)}$ x $2^{k(n)}$ Hermitian PSD matrix A
  - Promised either:
    - Minimum eigenvalue $\lambda_{min}(A) > b$ or $\lambda_{min}(A) < a$
    - Where $b-a > 2^{-O(k(n))}$
  - Which is the case?
- Compared to the **Local Hamiltonian** problem…
  - Input is *Succinctly Encoded* instead of Local
  - Precision needed to determine the promise is $1/2^k$ instead of 1/poly(n)
- *Our Result*: k(n)-**P.S Hamiltonian** problem is *complete* for **BQSPACE**[k(n)]

Upper bound (1/2):
k(n)-P.S Ham. $\in$ k(n)-bounded QMA$_{k(n)}$(c,c-$2^{-k(n)}$)

- *Recall*: k(n)-**Precise Succinct Hamiltonian** problem
  - Given Succinct Encoding of $2^{k(n)}$ x $2^{k(n)}$ Hermitian PSD matrix A, is $\lambda_{min}(A) \leq a$ or $\lambda_{min}(A) \geq b$ where $b-a \geq 2^{-O(k(n))}$?
- *Recall also*: Quantum algorithm for "phase estimation problem" [Kitaev '95]
  - Eigenvalues of unitary matrices are roots of unity, $e^{2\pi i\theta}$ for $0 \leq \theta < 1$
  - "Phase estimation problem": Given unitary U and eigenstate $|\psi\rangle$ output an approximation to the phase θ
- **PreciseQMA** protocol: Merlin sends eigenstate $|\psi\rangle$ with minimum eigenvalue
  - Arthur runs phase estimation with one ancilla qubit on e^{-iA} and $|\psi\rangle$

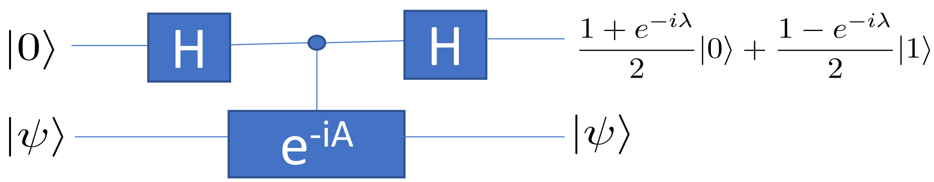

  - Measure ancilla and accept iff "0"
  - Easy to see that we get "0" outcome with probability that's slightly ($2^{-O(k)}$) higher if $\lambda_{min}(A) < a$ than if $\lambda_{min}(A) > b$
  - But this is exactly what's needed to establish the claimed bound!
- *Remaining question*: how do we implement e^{-iA} ?
  - We need to implement this operator with precision 2^{-k}, since otherwise the error in simulation overwhelms the gap!
  - Luckily, we can invoke recent "precise Hamiltonian simulation" results of [Childs et. al'14]
    - Given Succinct Encoding of A, implement e^{-iA} to within precision ε in space that scales with $\log(1/\varepsilon)$ and time $\text{polylog}(1/\varepsilon)$
- Using these results, can implement Arthur's circuit using O(k(n)) space and poly(n) time

Upper bound (2/2): $k(n)$-bounded $\mathbf{QMA}_{k(n)}(c, c-2^{-k(n)}) \subseteq \mathbf{BQSPACE}[k(n)]$

1. Error amplify the **PreciseQMA** protocol
   - *Goal*: Obtain a protocol with error inverse exponential in the witness length, $k(n)$
   - We want to do this while simultaneously preserving verifier space $O(k(n))$
   - We'll actually develop amplification technique that does this…
2. "Guess the witness"!
   - Consider this amplified verification protocol run on a *maximally mixed state* on $k(n)$ qubits
   - Not hard to see that this new "no witness" protocol has a "precise" gap of $O(2^{-k(n)})$!
3. Amplify again!
   - Use our "space-efficient" **QMA** error amplification technique again!
   - Obtain bounded error, at a cost of exponential time
   - But the space remains $O(k(n))$, establishing the **BQSPACE**[$k(n)$] upper bound

QMA amplification

- Our proof needed a particularly strong **QMA** amplification procedure
  - One that preserves both Merlin's witness length and Arthur's verification space
- Prior amplification methods
  1. "Repetition" [Kitaev '99]
     - Ask Merlin to send many copies of the original witness and run protocol on each one, take majority vote
     - *Problem with this*: number of witness qubits grows with improving error bounds
     - Needs $r/(c-s)^2$ repetitions to obtain error 2^{-r} by Chernoff bound
  2. "In-place" Amplification [Marriott and Watrous '04]
     - Define two projectors: $\Pi_0 = |0\rangle\langle 0|_{anc}$ and $\Pi_1 = V_x^\dagger |1\rangle\langle 1|_{out} V_x$
     - Notice that the max. acceptance probability of the verifier is maximal eigenvalue of $\Pi_0 \Pi_1 \Pi_0$
     - Procedure
       - Initialize a state consisting of Merlin's witness and all zero ancilla qubits
       - Alternatingly measure $\{\Pi_0, 1-\Pi_0\}$ and $\{\Pi_1, 1-\Pi_1\}$ many times
       - Use post processing to analyze results of measurements
     - Analysis relies on "Jordan's lemma"
       - Given two projectors, there's an orthogonal decomposition of the Hilbert space into 1 and 2-dimensional subspaces invariant under projectors
       - Basically allows verifier to repeat each measurement without "losing" Merlin's witness
       - Because application of these projectors "stays inside" 2D subspaces
     - As a result, we can attain the same type of error reduction as in repetition, without needing additional witness qubits

Space efficient In-place amplification

For other results improving Marriott-Watrous in various directions see e.g., [Nagaj et. al.'09 & F., Kobayashi, Lin, Morimae, Nishimura, **ICALP'16**]

- We're not happy with Marriott-Watrous amplification!!

$$k-\text{bounded } \mathsf{QMA}_m(c,s) \subseteq \left(k+\frac{r}{(c-s)^2}\right)-\text{bounded } \mathsf{QMA}_m(1-2^{-r},2^{-r})$$

  - The space grows because we need to keep track of each measurement outcome
  - We want to be able to *space-efficiently* amplify protocol with inverse exponentially small gap (i.e., c-s=$1/2^k$)
- We are able to improve this!

$$k-\text{bounded } \mathsf{QMA}_m(c,s) \subseteq \left(k+\log\frac{r}{c-s}\right)-\text{bounded } \mathsf{QMA}_m(1-2^{-r},2^{-r})$$

  - Now the same setting of parameters preserves O(k) space complexity!
  - Proof idea:
    - Define reflections $R_0 = 2\Pi_0 - I, R_1 = 2\Pi_1 - I$
    - Using Jordan's lemma:
      - Within 2D subspaces, the product R_0R_1 is a rotation by an angle related to acceptance probability of verifier V_x
    - Use phase estimation on R_0R_1 with Merlin's state $|\psi\rangle$ and ancillas set to 0
      - Key point: Phase estimation to precision *j* uses <u>*O(log(1/j))* ancilla qubits</u>
    - Accept if the phase is larger than fixed threshold, reject otherwise

Lower bound: k(n)-*Precise Succinct Hamiltonian* is **BQSPACE**[k(n)]-hard

- Follows from space-efficient **QMA** amplification and Kitaev's "clock-construction"
- Any language in **BQSPACE**[k(n)] can be decided by uniform family of quantum circuits $\{Q_x\}_{x\in\{0,1\}^*}$ of size at most $2^{k(|x|)}$
  - By our uniformity condition
- Kitaev shows how to take this circuit and build a Hamiltonian $H = \sum_{i=1}^{M} H_i$ with the property that:
  - In the "yes case", the Hamiltonian's minimum eigenvalue is less than some quantity a involving the *completeness* and the circuit size
  - In the "no case", the Hamiltonian's minimum eigenvalue is at least some quantity b involving the *soundness* and the circuit size
- By amplifying the completeness and soundness of the circuit we can ensure that the promise gap of the Hamiltonian, b-a, is at least 2^{-k}
- Easy to show that this Hamiltonian is succinctly encoded
  - Follows from sparsity of Kitaev's construction and uniformity of circuit

Application: **PreciseQMA=PSPACE**

- *Question:* How does the power of **QMA** scale with the completeness-soundness gap?
- *Recall*: **PreciseQMA**=$\cup_{c>0}$**QMA**($c,c-2^{-poly(n)}$)
- Both upper and lower bounds follow from our completeness result, together with **BQPSPACE**=**PSPACE** [Watrous'03]
  - *Upper bound* (**PreciseQMA**⊆**PSPACE**):
    - Showed poly(n)-**P.S Ham.** ⊆ **BQSPACE**[poly(n)]=**PSPACE**
  - *Lower bound* (**PSPACE**⊆**PreciseQMA**) :
    1. Showed poly(n)-**P.S. Ham.** is hard for **BQSPACE**[poly(n)]=**PSPACE**
    2. But also it's in **PreciseQMA** by "poor man's phase estimation"
- *Corollary*: "**precise k-Local Hamiltonian problem**" is **PSPACE**-complete
- *Extension:* "Perfect Completeness case": **QMA**($1,1-2^{-poly(n)}$)=**PSPACE**
  - *Corollary:* checking if a local Hamiltonian has zero ground state energy is **PSPACE-**complete

Where is this power coming from?

- Could **QMA**=**PreciseQMA**=**PSPACE**?
  - Unlikely since **QMA**=**PreciseQMA** $\Rightarrow$ **PSPACE**=**PP**
    - Using **QMA** $\subseteq$ **PP**
- How powerful is **PreciseMA**, the *classical analogue* of **PreciseQMA**?
  - *Crude upper bound*: **PreciseMA** $\subseteq$ **NP**$^{\mathbf{PP}}$ $\subseteq$ **PSPACE**
  - And believed to be strictly less powerful, unless the "Counting Hierarchy" collapses
- So the power of **PreciseQMA** seems to come from both the quantum witness and the small gap, together!

Understanding "**Precise**" complexity classes

- We can answer questions in the "precise" regime that we have no idea how to answer in the "bounded-error" regime
- *Example 1*: How powerful is **QMA(2)**?
  - **PreciseQMA**=**PSPACE** (our result)
  - **PreciseQMA(2)**=**NEXP** [Blier & Tapp'07]
  - So, **PreciseQMA(2)** ≠ **PreciseQMA**, unless **NEXP=PSPACE**
- *Example 2*: How powerful are quantum vs classical witnesses?
  - **PreciseQCMA** $\subseteq$ **NP**$^{\mathbf{PP}}$
  - So, **PreciseQMA** ≠ **PreciseQCMA**, unless **PSPACE** $\subseteq$ **NP**$^{\mathbf{PP}}$
- *Example 3*: How powerful is **QMA** with perfect completeness?
  - **PreciseQMA**=**PreciseQMA**$_{\mathbf{1}}$=**PSPACE**

Characterization 2:
Well-Conditioned Matrix Inversion

The Classical Complexity of **Matrix Inversion**

- The **Matrix Inversion** problem
  - Input: nonsingular n x n matrix A with integer entries, promised either:
    - $A^{-1}[0,0] > 2/3$ or
    - $A^{-1}[0,0] < 1/3$
  - Which is the case?

$$\mathbf{A} = \begin{pmatrix} a_{0,0} & a_{0,1}\cdots \\ \vdots & \\ a_{n,0} & a_{n,1}\cdots \end{pmatrix} \qquad \mathbf{A^{-1}} = \begin{pmatrix} ?\ldots & ? \\ \vdots & \\ ?\ldots & ? \end{pmatrix}$$

- This problem can be solved in classical $O(\log^2(n))$ space [Csansky'76]
- Not believed to be solvable classically in $O(\log(n))$ space
  - If it is, then **L**=**NL** (**Logspace** equivalent of **P**=**NP**)

Can we do better quantumly?

- "***Well-Conditioned* Matrix Inversion**" *can* be solved in *non-unitary* **BQSPACE**[log(n)]! [Ta-Shma'12] building on [HHL'08]
  - i.e., same problem with poly(n) upper bound on the condition number, κ, so that $\kappa^{-1}I \prec A \prec I$
  - *Appears* to attain quadratic speedup in space usage over classical algorithms
- *Begs the question*: how important is this "well-conditioned" restriction?
  - Can we also solve the *general* **Matrix Inversion** problem in quantum space O(log(n))?
  - Or could the Well-Conditioned case be in classical **Logspace**?

Our results on **Matrix Inversion**

- **Well-conditioned Matrix Inversion** is complete for *unitary* **BQSPACE**[log(n)]!
  1. We give a new quantum algorithm for **Well-conditioned Matrix Inversion** avoiding intermediate measurements
     - Combines techniques from [HHL'08] with amplitude amplification
  2. We also prove **BQSPACE**[log(n)] hardness– suggesting that "well-conditioned" constraint is *necessary* for quantum **Logspace** algorithms
- So this is another reason to believe Matrix Inversion can't be solved in classical **Logspace** (because otherwise **L**=**BQL**)

Can generalize from log(n) to k(n) qubits…

- *Result 3*: k(n)-***Well-conditioned Matrix Inversion*** *is complete for* **BQSPACE**[k(n)]
    - Input: Succinct Encoding of 2^k x 2^k PSD matrix A
        - Upper bound $\kappa<2^{O(k(n))}$ on the condition number so that $\kappa^{-1}I \preceq A \preceq I$
    - Promised either $|A^{-1}[0,0]| \geq 2/3$ or $\leq 1/3$
    - Decide which is the case?
- Additionally, by varying the dimension and the bound on the condition number, can use **Matrix Inversion** problem to *characterize* the power of quantum computation with simultaneously bounded time *and* space!

Open questions

- Can we use our **PreciseQMA**=**PSPACE** characterization to give a **PSPACE** upper bound for other complexity classes?
  - For example, **QMA**(2)?
- How powerful is **PreciseQIP**?
- Natural complete problems for *non-unitary* quantum space?

Thanks!