SAND2013-10370C

Implications of Analog Simulation for Computation and Complexity

Analog quantum simulation -- the kind of simulation you do in optical lattices, without error correction -- may or may not actually work. Oddly enough, everybody knows whether or not it will work... but half of us know that it will, and half know it won't! Both possible answers have interesting implications for computational complexity. In addition to discussing these implications, I'll use analog simulation as inspiration to try and answer the question "What kind of algorithms could run usefully on a small quantum computer?"

Robin Blume-Kohout

(Sandia National Labs)

U.S. DEPARTMENT OF ENERGY

Sandia National Laboratories

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Overview

- We're all wondering "*What do we do with a small quantum computer?*".
- I'm particularly curious about "What can a small QC do?" This is kind of like making stone soup: we have very little quantum memory, and we can't do error correction. So where should we look to find awesome?
    - *1st*: **Potentially useful small-QC algorithms.**
    - *2nd*: **Minimalist error correction.**
    - *3rd*: **Why analog simulation is provocative.**
    - *4th*: **Seeking other intrinsically robust algorithms.**

I. A Hasty Look at Small-QC Algorithms

QLOGSPACE

- One of the defining features of "small quantum computer" is that it doesn't have very many qubits.
- I would like to solve problems about $N>>1$ classical bits.
- O(1) qubits ==> finite state automaton = pretty limited.
- So let's let the QC have O(log N) qubits. The problems it can solve are in **QLOGSPACE** (by definition).
- **Note 1**: Classical LOGSPACE actually includes some fairly interesting things! All of arithmetic, and pretty much the whole class of streaming algorithms.
- **Note 2**: This motivates looking for *space* efficiency, rather than *time* efficiency. Which is good, because O(1) qubits can probably be simulated classically...

Le Gall's algorithm

- So can a quantum computer do anything in less space than a classical machine? (Is anything in **QLOGSPACE** but not **LOGSPACE**?)
- Yes! Le Gall used a Grover-based quadratic separation in communication complexity to show...

Theory of Computing Systems
August 2009, Volume 45, Issue 2, pp 188-202

Exponential Separation of Quantum and Classical Online Space Complexity

François Le Gall

✦**Input:** Two m-bit strings $\boldsymbol{x}$ and $\boldsymbol{y}$ are fed in, $m^{1/2}$ times in a row.

✦**Problem**: Is there any index i for which $x_i = y_i = 1$?

✦**Classically**: The computer must have at least $m^{1/2}$ bits of memory!

✦**Quantumly**: Problem can be solved in only $\log(m)$ qubits of memory!

Useful streaming algorithms?

- Le Gall's algorithm is frankly pretty contrived:

    1. When are you going to get the same two strings streamed through *successively* that many times? (Note: if they come in parallel, the problem becomes trivially easy).

    2. For the QC to be *useful*, we need $\sqrt{N}$ classical bits to be expensive. Which means $N >> 10^{12}$... and now our QC needs 10^6 repetitions of that terabit string! That's going to take absurdly long just to read.

- Open question: Are there *real* problems for which QC allows significant memory reduction?

    - Search space: known big/streaming data problems where memory is the bottleneck.
    - Example: In a stream of N letters, are any repeated? Requires O(N) memory.
    - For these sorts of problems, even an $N \dashrightarrow \sqrt{N}$ reduction in memory could be significant, and an exponential reduction would be huge.

II. Error Correction (Grudgingly)

We need error correction

- Suppose we found such an algorithm. Imagine it does something awesome -- analyzes N=2^{30} classical bits -- using only 30 qubits.
- But we still have to read in N bits. 1 trillion timesteps. And each timestep, we have to perform at least one gate.
- If $p_{\text{fail}} = 10^{-9}$, our device *still* dissolves into useless mush!
- Conclusion: If doing something useful with 30 qubits requires reading a lot of data, then by the time we read in all the data, decoherence will ruin everything unless we do some kind of error correction.

Error correction is expensive

- If we use standard quantum error correction (QEC), and we only have 30-100 physical qubits...
  ...we aren't going to have enough logical qubits to do *anything* interesting.

- Fault tolerant QEC provides a protected logical cocoon, within which *any* encoded algorithm will work. This is great -- but it comes at too high a cost.

- My central question for today: **are there less expensive ways to protect *particular* algorithms against errors?**

Entangling EC & Algorithms

- If we proceed this way, we can't separate algorithms and error correction.
- We can *start* by asking "What is essential for EC?", but if we stay agnostic about algorithms, then the answer is ultimately going to be "The full, expensive machinery and overhead of fault tolerance."
- So, ultimately, this is about a weaker notion of fault tolerance -- or, alternatively, about algorithms that are *intrinsically robust* to noise, and provide some or all of their own fault tolerance.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

The essence of EC: Cooling

- Suppose we apply a succession of *unital* operations to our small QC -- and each one involves a little depolarization.
- We inexorably end up in the maximally mixed state.
- So *non-unitality* is the critical resource -- a.k.a. cooling.
  - Without cooling, we can't keep information alive.
  - Given some cooling + near-unitary control, we can distill pure ancillas and use them to do FT QEC.
- In fact, *any* error correction circuit can be viewed as nothing but cooling (in a very weird decomposition!)

Is EC nothing but cooling?

- Any error correction process can be transformed into pure cooling of a "syndrome" subsystem. However...
- ...this obfuscates a critical point: the subsystem being cooled is *very* weird and unnatural.
  - highly nonlocal
  - demands extreme precision
  - computationally hard to implement cooling.
- So in one sense, error correction is nothing but cooling. But it's not "dumb cooling" -- it's a very smart, clever, *contrived* cooling process.

Self-correcting systems

- Wouldn't it be nice if we had a "natural" fault tolerant QEC scheme, where simple physical cooling was enough?
- This is the definition of a *self-correcting quantum memory*. The critical ingredient is a free energy landscape that surrounds the code states with a broad funnel.
- Self correcting memories aren't known to exist in less than 4 spatial dimensions, and even if they do, they're probably hard to build... and will still require a lot of extra qubits.
- So I'm going to assume for now that a self-correcting fault tolerant universal QC is waaayyyy too hard.

Self-correcting *algorithms*

- Something that might be easier: Can *specific* algorithms be protected against errors just by [local, physical] cooling?
- This might be easier than protecting *arbitrary* algorithms.
- This is the point where we have to start considering specific algorithms, implemented in specific ways, and their interaction with cooling / nonunitality / dissipation.
- Natural first question:
  "Are there any known algorithms that can run effectively without full-scale FT/QEC?"

III. The Promise of Analog Simulation

Why analog simulation

- Because it's is the only algorithm I know of that even *hints* at this kind of intrinsic robustness.
- So I want to examine analog simulation, and its potential implications for computational complexity.
- Do not expect rigor, math, or concrete results! Even this core section is primarily "ideas"-focused.
- It's okay if this makes you angry. Go prove that I'm wrong about everything! (That would be a cool result).

Quantum Simulation

...other stuff...

Factor!
(Shor)

Quantum Computer
(What can it do?)

Search!
(Grover)

Simulate Quantum Systems!
(Feynman/Lloyd/Zalka)

Tell you anything a perfect experiment could = "**weak simulation**"

Calculate wavefunction to reasonable *relative* precision = "**strong simulation**"

U.S. DEPARTMENT OF ENERGY
Sandia National Laboratories

Digital weak simulation

- Quantum computers can definitely do weak simulation.
- If we happen to have a fault tolerant QC sitting around, we can implement $U = e^{-iHt}$ for any reasonable H by simply Trotterizing it (see Lloyd 1996).
- So given any $|\psi(0)\rangle$, we can produce $|\psi(t)\rangle = U\,|\psi(0)\rangle$ and then measure any desired property of it.
- This presumes that we can make the initial state, and that by "simulate" we mean "simulate time evolution".

Dynamic vs. Static properties

- This brings up another great divide:

People who care about e^{-iHt} **People who care about $|0\rangle$**

- I'm going to focus on dynamics.
- However, it's not clear that these are different questions.
- It also brings up a question that confused me for a long time:
  "**Why are condensed matter physicists obsessed with the ground state?**"

Do we still need EC?

- So if we had a full-on QC, we could do simulation.
- But do we need error correction and fault tolerance for it?
- Quantum simulation isn't just any quantum algorithm. Nature is *really* good at implementing $e^{-i\boldsymbol{H}t}$!
- So instead of painstakingly implementing Trotterization, how about we just *build* $\boldsymbol{H}$? This is "analog simulation".

Analog[ue] Simulation

- Why "analog"?

  - Analog classical computing uses continuous variables and *cannot* be error corrected.

  - Analog quantum simulation can still be discrete, but *doesn't* use error correction. (Could it be???)

- Arguably, "analogue" is a better term:

  - We are building a system that serves as an analogue or proxy for the one we want to study (tunable graphene, Fermi-Hubbard model, whatever...).

  - We can measure any property of interest on the analogue, and (if it works), it will mimic the original system.

U.S. DEPARTMENT OF ENERGY
Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Does digital dominate analog?

- We know that analog classical computing is <= digital classical computing in the presence of any noise at all.
  (but before digital was available, analog was useful!)

- By similar arguments, analog[ue] quantum computing is surely <= digital quantum computing (thanks to Trotterization, we can simulate A with D).
  (but digital QC is not available yet...)

- So in the absence of powerful digital quantum computers, there is room for analog *quantum* computing to solve problems that digital *classical* computers can't.

U.S. DEPARTMENT OF ENERGY
Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Analog, digital, analogue...

	Analog	Digital
Analogue	wind tunnel LHC	Quantum analog[ue] simulations of graphene, Hubbard, Ising, etc.
Algorithmic	analog classical computers (FAIL)	Shor, Grover, FFT, PageRank... every other algorithm

Can analogue simulators work?

- "**YES**." Why? It's all about the Hamiltonian. Same H, same system. Similar H, similar system. So if we implement the right Hamiltonian (approximately), we'll get the right behavior.

- "**NO**." Why? Errors will propagate. As soon as the system's state evolves into one that's orthogonal to the "correct" one, all bets are off.

- **MY BELIEF**: Yes. Why?
  1. We don't care about micro-behavior.
  2. Macro-variables that are sensitive to perturbations are random in practice -- so we ignore them! We only care about simulating *robust* properties.
  3. Ergo, quantum simulators with small errors will still work.

U.S. DEPARTMENT OF ENERGY
Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

“Analogizing” Perturbations

- OBJECTION: “Property X may be stable under lab perturbations, but not under analogizing perturbations that occur when we map $H_{graphene}$ --> $H_{optical\ lattice}$.”

- REPLY: A valid concern (which I can’t put to rest). But recall that thermodynamic quantities are analytic w/r.t. changing thermodynamic parameters. So as long as Property X is a thermodynamic quantity, it should vary smoothly...

  ...except that “analogizing” perturbations are not variations in thermodynamic parameters. So we have no guarantee that robust-in-the-lab parameters are necessarily robust-under-analogization. So we should try to prove/disprove this!

Is AQSIM in BPP?

- OBJECTION: "Maybe analog quantum simulators will work for [some] robust properties -- but only ones that can be calculated efficiently by *classical* computers too!"
- REPLY #1: Why on earth would you say that???
- REPLY #2: So what about the conductivity of the Fermi-Hubbard model? Is that (a) not robust under analogization (why?) or (b) in BPP (why haven't we solved it?).
- If F-H simulations *don't* work, that would be cool -- there would be new physics about analogizing that we don't understand. Why do wind tunnels work, but not simulations of the F-H model?

Complexity Implications

- If analog quantum simulators *do* work, then either:

1. There is a BPP algorithm for simulating F-H (and every other "robust" Hamiltonian).

$\Downarrow$

Surprising! but boring.

2. An analogue simulator can calculate something ***not*** in BPP.

$\Downarrow$

Not totally shocking (still in BQP, so doesn't imply analogue > digital

...

But... we believe QCs without clever error correction are *not* BQP-complete (at least in 2D).

$\Downarrow$

New complexity class AQSim(2D)? (Contained in BQP but not BPP).

Error Correction Implications

- Case 2 (the interesting one) also implies some interesting things about error correction and fault tolerance.
- Recall: "analogue" = "no clever EC", not "no EC at all"!
- With no EC, any system depolarizes very rapidly. A simulator must be maintained at low temperature (cooled).

Conventional view of EC:
- ✦ codes (linear/stabilizer)
- ✦ syndrome extraction
- ✦ decoding
- ✦ correction

"Liberal" view of EC:
- ✦ pump entropy out
- ✦ cooling (usually clever)

- So an analog simulator does use and require EC... but just dumb cooling. And this is sufficient for fault tolerance.

Is "dumb cooling" enough?

- A computer for which local, physical cooling enables robustness is a self-correcting memory (and more -- FT!).
- This is hard quantumly, but around 300 K, practically everything around us is a self-correcting *classical* memory (rocks, tables, stars, ... ants, humans...)
- Furthermore, many of those things aren't just memories. **They also compute their own dynamics robustly** (and repeatably -- they really are computing something).
- This is pretty mundane for classical objects (we only get excited about *universal* computing). But for a *quantum* object, computing its own dynamics (efficiently) might defeat classical universal computers!

Intrinsic Robustness

- Suppose graphene, the Fermi-Hubbard model, etc. are:
  (1) hard to simulate classically,
  (2) capable of reliably computing their own dynamics.
- Then they are natural, fault-tolerant, 2D "computers" that solve a problem outside of BPP.
- They're certainly not BQP-complete, but they do *something* superclassical -- and they do it robustly without "clever" error correction and fault tolerance.
- Even if analogue simulations of these models don't work, the existence of intrinsic robustness is interesting!

U.S. DEPARTMENT OF ENERGY
Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Summary Conjecture

- Let me conclude this discussion with a conjecture that summarizes most of it and states *my* best guess:

  There exists a well-defined quantum system that:
  (1) has macroscopic dynamical properties that are robust in the lab;
  (2) these properties are robustly computable by analog quantum simulators at low but nonzero temperature;
  (3) these properties cannot be computed in BPP.
  This system defines a novel complexity class AQSim.

- Note: In order to really place this in the context of complexity, the system must be *parameterized* -- we need an input to the problem, which should be specified by the parameters of the system simulated.

IV. Beyond Simulation

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Simulation isn't enough

- Quantum computers will almost certainly be good at quantum simulation. It's a good trick... but limited.
- I think simulating physical systems is pretty interesting. So do many physicists, chemists, and other scientists.
- However, it's a tiny fraction of the *economic* value of computation -- and therefore not as compelling to funding agencies as we might wish.
- So let's address a different question:
  Can we devise more traditional "algorithms" that inherit the [possible] virtues of analog simulation -- i.e., super-classical performance and intrinsic robustness to noise?

An Ambitious Goal

- GOAL: Algorithms for [small] quantum computers that:

(1) Are clearly digital,
(2) Process [large amounts of] classical data,
(3) Run for a pretty long time,
(4) Provide well-defined answers to [decision] problems,

and

(5) Do not require clever error correction to do all that.

U.S. DEPARTMENT OF ENERGY
Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Mechanisms

- If we had such an algorithm (I don't)... how might it work?
- Clearly we need entropy reduction ("liberal EC"). Cooling is within the rules. Most conservative approach is just contact with a cold bath.
- But we can go further: allow dissipative operations (e.g. intentional amplitude damping) as part of a gate set. Assume pretty high fidelity (comparable to unitary gates).
- One appealing model is dissipative computing, in which the "gates" are partly or exclusively nonunitary and nonunital operations.
- PUNCHLINE: If universal FT/EC is too hard, fall back on what we *know* we need -- entropy reduction -- incorporate it into the gates themselves, and push computation forward using non-unital operations (somewhat reminiscent of measurement-based QC).

Algorithms

- In traditional QC, error correction and computation are separate. The purpose of EC is to prepare a cocoon in which ideal logical computation is simulated.

- Intrinsic robustness couples computation and EC tightly -- which constrains the possible algorithms. To what?
  (1) I don't know!
  (2) Prepare low-temperature thermal states? (boring...)
  (3) Encode computational problems into Hamiltonians?
  Lots of research on this, but most seems to be about *proof* classes, not forward computation (history state AQC is a major exception).

Phase Transitions

- I have a hunch:
  Intrinsic robustness $\Leftrightarrow$ 2+ phases, separated by transitions.
- These don't have to be thermal/quantum phases of a Hamiltonian (consider, e.g. stable behavior of dissipative master equations... what else could have a phase?)
- Answer to a decision problem is encoded in which phase we end up in -- which, of course, is robust to perturbations *and* can be read out easily.
- Input instance is used to define the system (e.g. Hamiltonian) in which the phases appear.
- Can we say anything generally about this model -- e.g., prove that certain problems *can't* be mapped to phase transitions this way?

U.S. DEPARTMENT OF ENERGY Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Conclusions

- Small quantum computers are at a disadvantage for computing in less *time*, but they are well-suited for computing things in less *memory.*
- However, this requires reading long inputs...
  ...which seems to require some form of error correction.
- Cooling might be enough to make *some* algorithms robust. Analog simulation is a potential example of such an algorithm.
- Analog simulation -- if it works -- defines an interesting complexity class.
- If intrinsic robustness can be exported to more "algorithmic" problems, doing so may require mapping those problems to systems with multiple phases and phase transitions between them.

U.S. DEPARTMENT OF ENERGY Sandia National Laboratories
Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.