Active FTQEC—The Curse of the Open System

Amit Hagar

Department of History & Philosophy of Science
College of Arts and Sciences
Indiana University, Bloomington

Supported by NSF Scholar's Award # SES–0847547

CRPC Workshop
Oxford Computing Lab
Aug 25, 2009

Motivation

Why there is still no large scale, fault tolerant and computationally superior quantum computer?

- Optimists: It's only a technological problem, and we still lack the technology.

  Call this thesis (1)

- Pessimists: Irrespective of technological capabilities, there are *in principle* obstacles.

Motivation

If you believe quantum mechanics is going to break down before nontrivial quantum computing becomes possible, then you must believe theres some point where it will break down—some level of size, or complexity, or whatever, at which it will cease to be a useful description of the world.

Aaronson 2004

Motivation

Which is just another way of saying that if you are a pessimist, then you must also deny the universal applicability of QM.

Call this thesis (2).

If you believe quantum mechanics is going to break down before nontrivial quantum computing becomes possible, then you must believe theres some point where it will break down—some level of size, or complexity, or whatever, at which it will cease to be a useful description of the world.

Aaronson 2004

Motivation

But why are theses (1) & (2) so widely embraced?

- There are, as always, sociological factors (personality issues, inter–disciplinary issues).
- And, more seriously, there are the threshold theorems.

Therefore, noise, if it is below a certain level, is not an obstacle to unlimited resilient quantum computation.

Knill *et al.* 1998

Motivation

- The threshold theorem tells us that, in principle, we will be able to construct devices to perform arbitrarily long quantum computations using a polynomial amount of resources, so long as we can build components such that the per–gate error is below a fixed threshold. In other words, noise and imprecision of physical devices should not pose a fundamental obstacle to realizing large–scale quantum computers . . . the theorem has given confidence that they can be built.

  Kaye *et al.* 2007

- The [threshold] theorem made it clear that no physical law stands in the way of building a quantum computer.

  Gaitan 2008

Motivation

- The existence of fault-tolerant schemes turns the problem of building a quantum computer into a hard but possible-in-principle engineering problem: if we just manage to store our qubits and operate upon them in a level of noise below the fault-tolerance threshold, then we can perform arbitrary long quantum computations.

  Kempe, Regev, Unger, and Wolf, 2008

Motivation

- The purpose of this talk is to raise doubts with respect to thesis (2) (and subsequently with respect to thesis (1)) by questioning the physical significance of the threshold theorems.
- This will be done by summoning physics (Thermodynamics and statistical mechanics) & HPS (the foundations thereof).
- On the positive side, I'll end with a hint of how I think skepticism w.r.t large scale, fault tolerant and computationally superior QC could be made precise.

Outline

1. Background
   - Error correction
   - Fault Tolerance - the Active Approach
2. Pessimism
   - Physics
   - History & Philosophy of Science
3. Current Situation in Active FTQEC
   - Non Markovian Noise Models
4. Morals
   - Optimal Skepticism
   - Let's Go Passive

Outline

1. Background
   - Error correction
   - Fault Tolerance - the Active Approach
2. Pessimism
   - Physics
   - History & Philosophy of Science
3. Current Situation in Active FTQEC
   - Non Markovian Noise Models
4. Morals
   - Optimal Skepticism
   - Let's Go Passive

Outline

1. Background
   - Error correction
   - Fault Tolerance - the Active Approach
2. Pessimism
   - Physics
   - History & Philosophy of Science
3. Current Situation in Active FTQEC
   - Non Markovian Noise Models
4. Morals
   - Optimal Skepticism
   - Let's Go Passive

Outline

1. Background
   - Error correction
   - Fault Tolerance - the Active Approach
2. Pessimism
   - Physics
   - History & Philosophy of Science
3. Current Situation in Active FTQEC
   - Non Markovian Noise Models
4. Morals
   - Optimal Skepticism
   - Let's Go Passive

Surpassing the No Cloning Theorem

Suppose we would like to protect a qubit against a bit–flip. We encode it into a larger space *without* copying or measuring it:

$$\alpha|0\rangle + \beta|1\rangle \rightarrow \alpha|000\rangle + \beta|111\rangle \tag{1}$$

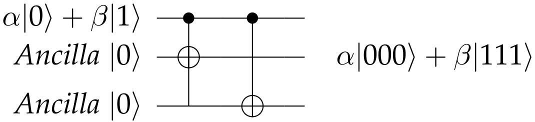

(2)

If a bit–flip happens now, and the superposition in Eq. (1) becomes, say, $\alpha|100\rangle + \beta|011\rangle$, we can still extract information from the state without destroying it by measuring the *parity* of all pairs of qubits.

Surpassing the No Cloning Theorem

For instance we can measure the parity of the first two qubits with the following circuit

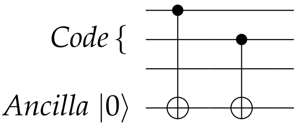

(3)

Here each CNOT flips the ancilla qubit if the source qubit is in the state $|1\rangle$. If the first two qubits are in the state $|00\rangle$, the ancilla is left in the state $|0\rangle$. If these qubits are in the state $|11\rangle$ the ancilla is flipped twice and returns to state $|0\rangle$. Otherwise it is flipped once by one of the CNOTs.

Surpassing the No Cloning Theorem

A more complicated encoding exists for phase–flip errors, that uses 9 qubits and can also correct a bit–flip error and a combination of both:

$$|0\rangle_{enc} = \frac{1}{\sqrt{2^3}}(|000\rangle + |111\rangle)(|000\rangle + |111\rangle)(|000\rangle + |111\rangle) \quad (4)$$

$$|1\rangle_{enc} = \frac{1}{\sqrt{2^3}}(|000\rangle - |111\rangle)(|000\rangle - |111\rangle)(|000\rangle - |111\rangle) \quad (5)$$

With this encoding each of the blocks of three qubits is still encoded with a repetition code, so we can still correct bit flip errors in a fashion very similar to the above.

Surpassing the No Cloning Theorem

To detect a phase flip without measuring the information in the state we use Hadamard gates to change bases from the standard basis to the $|\pm\rangle$ basis

$$|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \;\; ; \;\; |-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \tag{6}$$

and measure the parity of the phases on each pair of two of the three blocks in the new basis (a phase flip in the standard basis becomes a bit flip in the $|\pm\rangle$ basis).

Discretizing the Errors

- Any unitary transformation that the composed system may undergo can be expressed as a linear combination of four basic errors (the *Pauli group*).
- By performing a ‘syndrom’ measurement we collapse the combined state on one of the four ‘error subspaces’ hence disentangle the error from the information stored in the qubit without destroying it.
- Depending on the outcome of the syndrome measurement, we can correct the error by applying the respective operation to the appropriate qubit.

What Do we Gain?

The redundancy allows one to improve on the error probability for a single qubit ϵ:

- The state will be projected onto either a state where no error has occurred with probability $1 - 9\epsilon$, or onto a state with a large error (single qubit, two qubit *etc.*).
- Such a code protects against all single qubit errors. Only when two (independent) errors occur (which in this case happens with probability $\leq 36\epsilon^2$), the error is irrecoverable.
- Thus Shor's 9 qubit QEC is advantageous whenever $\epsilon \leq 1/36$.

The Challenge

- With large number of errors t per code block, one reaches a point where the error–recovery procedure takes too much time, that it becomes likely that $t + 1$ errors occur in a block, and the error–correction would fail.
- To keep this failure probability much smaller than 1, the error rate ϵ must *decrease* with the length of the computation; the longer the computation, the more accuracy it requires.

The Challenge

- The number of steps required for recovery scales as a power of t, t^a with exponent $a > 1$.
- The probability to have $t + 1$ errors before a recovery step is completed scales as $(t^a \epsilon)^{t+1}$.
- This expression is minimized when $t = c\epsilon^{-(1/a)}$ for some constant c and its value is at least $p = \exp(-ca\epsilon^{-(1/a)})$.
- The probability to fail per error correction cycle is at least p. If we have N such cycles, our total failure probability is $Np = \exp(-ca \log N\epsilon^{-(1/a)})$.
- For $p \ll 1$, ϵ must scale as $(1/\log N)^a$.

The Miracle

Concatenated codes involve recursively re–encoding already encoded bits. In the first–level one encodes each qubit with an appropriate code. Then, for each of the codewords one encodes each of the qubits again using the same code.

- $N \rightarrow Npoly(\log N)$.
- error probability per gate p reduces from cp^2 to $c(cp^2)^2 = c^3p^4$ for some constant c that depends on the code.
- This improves the error rate exponentially as long as $p < 1/c$.
- If one uses k levels of concatenation, the error at the highest level is reduced to $\frac{(cp)^{2^k}}{c}$.

The Miracle

- The error rate decreases faster than the growth in the size of the circuit.
- An error threshold exists such that if each gate in a physical implementation of a quantum network has error less than this threshold, it is possible to perform an arbitrary long quantum computation with arbitrary accuracy.

Assumptions of active FTQEC

1. Error correlations decay exponentially in time and space.
2. Gates can be executed in time τ_g such that $\tau_g \omega = O(\pi)$, where ω is the Bohr or the Rabi frequency.
3. A constant supply of ‘fresh’, nearly pure, ancilla qubits is available.

Derivation of the Markovian Master Equation (MME)

- Two types of fully rigorous derivations of quantum MME: the singular coupling limit (SCL) and the weak coupling limit (WCL).
- Both derivations must satisfy a thermodynamic constraint (KMS), namely that the reservoir is in a state of thermal equilibrium.
- Within the SCL, this condition allows the reservoir's correlation function to be approximated by a delta function only in the limit $T_R \to \infty$ where T_R is the reservoir's temperature.
- Physically, 'zero–memory', or assumption (1), holds within the SCL only when the reservoir is *much* hotter than the system on the same energy scale set by the system + ancillas.

Derivation of the Markovian Master Equation (MME)

- The final point, while consistent with assumption (2)—allowing arbitrary gate velocities—contradicts assumption (3): where does one get 'fresh, almost pure' ancillas to dump entropy in if $T_R \gg T_S$?
- Inversely, if one requires pure ancillas, then by coupling them to the system one must abandon the Markovian noise model in the environment.

Derivation of the Markovian Master Equation (MME)

The more realistic domain of WCL appears to be as unfavorable to FTQEC as the SCL:

- Within the WCL (where T_R is finite), one can achieve the Markovian condition in the reservoir's correlations function only after coarse graining over very long time–scales.

Expressed in terms of the system's gates frequency, this condition, while consistent with assumption (3), violates assumption (2), as it only allows slow, adiabatic, gates.

Derivation of the Markovian Master Equation (MME)

To summarize:

- (1) and (2) are incompatible with WCL, and so require SCL, which means that the reservoir (the source for the ancillas) must posses a high temperature, which then contradicts (3).
- (1) and (3) are incompatible with SCL, and so require WCL, which means that the gate velocity must be slow, which then contradicts (2).

Alicki, Lidar & Zanardi 2006

Those who cannot remember the past are condemned to repeat it

Note the timeline:

- 1994: Shor's Algorithm
- 1995: QEC
- 1996-7: FTQEC and the threshold theorems (TT) for Markovian noise
- 2006: First published criticism on the physical significance of TT with Markovian noise.

The following historical detour into the foundations of TD and SM might have saved the quantum information industry some time (and the funding agencies some money)

Two Problems in the Foundations of SM

- IRR: Explaining Irreversibility (or the TD arrow in time)
- PROB: Justifying Probability (the uniform measure and other probabilistic assumptions)

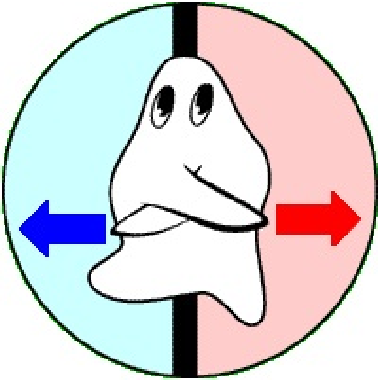

The Open System Approach

Interventionism is an attempt to answer IRR, and it comes in two flavors:

- Benign: Random noise kicks the system out of those states that lead to abnormal TD behavior and sets it back on a normal TD course.
- Radical: Eliminate the noise ("close thy system") and no thermalization would take place.

Note that as a solution to PROB, the benign version is also questionable: what justifies the double standard with respect to the dynamics of the system and those of the environment?

Analogies

- QC: an open system out of thermodynamic equilibrium.
  Noise: arising from computer's interaction with a heat bath.
- Active QEC: the attempt to "cool down" the open system, thus preventing its thermalization.
  TT: the promise that given a certain noise–level, this "prevention" can be done for an *arbitrarily long* time without increasing the *overall* thermodynamic cost.
- *Had the latter increased, we would have returned to (classical) irreversible computation, contrary to the presumably unitary (hence reversible) quantum computation that is taking place.*

1st Reason for Suspicion

The latent double standard with which the system and the environment are treated:

- While the interaction between the qubits + ancillas is entangling (and so non–local correlations dynamically evolve), in the interaction with the environment (or the noise model) no non–local correlations are allowed to evolve.
- With every computational step the environment acts as if it has "seen" the system for the first time.
- Given this double standard, the original TT seem now less miraculous: if one is allowed to cheat just once in quantum mechanics, one can indeed do miracles.

2nd Reason for Suspicion

If the composed system were left to itself, it would eventually equilibrate, wouldn't it?

- Radical interventionism: it is the noise that is responsible for the thermalization of the quantum computer; if we eliminate the noise, no such process would take place.
- But TD tells us that *all* physical systems out of equilibrium thermalize, and SM only changes the "all" to "almost all".
- External perturbations (like stirring a bowl of hot soup) may accelerate this process, but, apart from radical interventionists, no one sees them as necessary for thermalization.

Choices and Their Consequences

So FTQEC must make a choice: is it Radical or Benign?

- If Radical, then the history of science tells us it is suspect.
- If Benign, then noise–elimination is not the focus. Rather it is the abnormal (noise–resilient) states that we should be trying to create.

Non–local Noise

- Local noise: acting on the qubits *independently* of the structure of the evolution of the quantum computer. This evolution only *propagates* the errors.
- Non–Markovianity: the environment now “sees” the evolution of the computer and “learns” it. Since this evolution is non–local, it will eventually give rise to non–local noise.

Non–local Noise

- Although the strength of this effect can be mitigated by lowering the velocity of the quantum gates (i.e., by increasing the overall computational time), its existence is a *generic* consequence of any interaction and cannot be eliminated.
- The unavoidable interaction with the vacuum already introduces long–range quantum memory which causes the environment to be rather malevolent by tracing the (necessarily entangled) evolution of the system.
- The more entangling is the evolution of the quantum computer, the more non–local is the noise.

Its effect on active FTQEC

This question is still open, but:

- Length of the quantum computation? requires a delicate analysis of different time scales.
- Error–rates? TT must now deal with *amplitudes* and not with *probabilities*. Thresholds now become much lower than the previous (uncorrelated) case.
- TT for correlated (non–Markovian) noise explicitly rely on the norm of the interaction Hamiltonian. Low error–rate $\equiv$ very–high–frequency component of the noise is particularly weak.
- Not physically well motivated: in some decoherence models it even implies that the system and the environment are practically *de*coupled.

Upshot

The initial optimism that followed the discovery of QEC and active FTQEC seems now a little premature.

- FTQEC is not only contingent upon technology but also dependent on the actual noise model. For some noise models FTQEC might be impossible, while for other it may still be within reach for a certain amount of time.
- The irony is that the attempt to characterize *actual* quantum noise requires exponential resources. It seems that we need a quantum computer to tell us whether a large–scale, fault–tolerant and computationally superior quantum computer is possible.

It's the noise model, stupid

- Pessimism with respect to the feasibility of large–scale, fault–tolerant and computationally superior quantum computers is far from *ideological*.
- One need not abandon quantum theory in order to doubt the existence of such machines, and, on the other hand, the obstacles in realizing such machines need not be deemed purely technological.

It's the noise model, stupid

- A new type of skepticism is required: one which is not too strong (as it acknowledges the universal applicability of quantum theory), and at the same time is not too weak (as it isn't continegt upon technological capabilities).
- Here I have only argued that such a skepticism is *possible*; it is up for the pessimists to make it precise.

The Subsystem Approach

- Since the distinction between the system and the noise is completely arbitrary from a fundamental perspective, active FTQEC seems to be the wrong way to approach the project.
- In fact, *active error correction* seems to be a misnomer; the essence of the project should be *passive* rather than *active*. In other words, errors should be avoided and not corrected.
- For more on this passive view, see the Subsystem Approach (Zanardi, Knill, Viola, Lidar,. . .)

Bonus: a Hint For the Pessimists

Suppose noise–resilient and computationally superior states are "rare". What does it mean operationally?

- The resources for creating the them scale exponentially.

or

- The subspace in which these states reside scales down exponentially relative to the entire state space.

or

- The measurement rate for keeping those states in their resilient subspace scales exponentially.

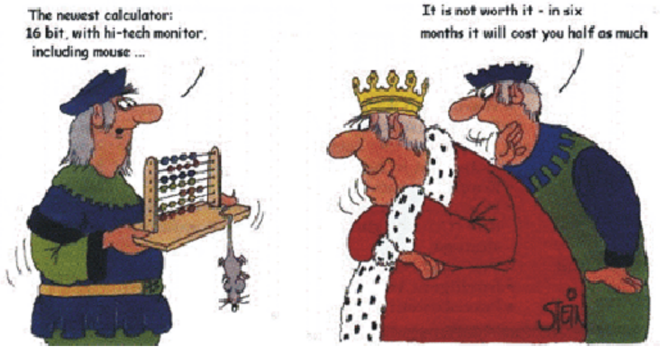
The newest calculator:
16 bit, with hi-tech monitor,
including mouse ...
It is not worth it - in six
months it will cost you half as much
STEIN