Active Fault-Tolerant Quantum Error Correction: The Curse of the Open System*

Amit Hagar†‡

Relying on the universality of quantum mechanics and on recent results known as the "threshold theorems," quantum information scientists deem the question of the feasibility of large-scale, fault-tolerant, and computationally superior quantum computers as purely technological. Reconstructing this question in statistical mechanical terms, this article suggests otherwise by questioning the physical significance of the threshold theorems. The skepticism it advances is neither too strong (hence is consistent with the universality of quantum mechanics) nor too weak (hence is independent of technological contingencies).

1. Introduction. A computer is a physical object; as such it is subject to the laws of physics. No doubt the application of these laws is error free, but external influences and hardware imprecisions can sometimes cause a mismatch between what the computer is supposed to do and what it actually does. In quantum computing, where algorithms exist that can solve computational problems more efficiently than any known classical algorithms,[1] the elimination of errors that result from external disturbances or from imperfect gates has become a "holy grail," and a worldwide

*Received June 2009; revised August 2009.

†To contact the author, please write to: Department of History and Philosophy of Science, College of Arts and Sciences, Indiana University, Bloomington, IN 47405; e-mail: hagara@indiana.edu.

‡This research is supported by the NSF (Grant SES no. 0847547) and Indiana University's Office of the Vice President for International Affairs. Any opinions, conclusions, or recommendations expressed in this material are mine and do not necessarily reflect the views of the NSF. I thank Itamar Pitowsky for helpful comments and Gerardo Ortiz and Osvaldo Pessoa for discussion. I also thank two anonymous referees for helpful suggestions and comments. I am solely responsible for any mistakes or errors that might appear herein.

1. For an accessible review, see Aharonov 1998.

Philosophy of Science, 76 (October 2009) pp. 506–535. 0031-8248/2009/7604-0005$10.00
Copyright 2009 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

quest for a large-scale, fault-tolerant, and computationally superior quantum computer is currently taking place. Whether such machines are possible is an exciting open question, yet the debate on their feasibility has been so far rather ideological in character (Wolfram 2001; Levin 2003; Goldreich 2005; Dyakonov 2006). Remarkably, philosophers of science have been mostly silent about it: common wisdom has it that philosophy should not intervene in what appears to be (and is also presented as) an engineering problem, and, besides, the mathematics employed in the theory of fault-tolerant quantum error correction (FTQEC) is rather daunting. It turns out, however, that behind this technical veil the central issues at the heart of the debate are worthy of philosophical analysis and, moreover, bear strong similarities to the conceptual problems that have been saturating a field quite familiar to philosophers, namely, the foundations of statistical mechanics (SM). Reconstructing the debate on FTQEC with statistical mechanical analogies, this article aims to introduce it to philosophers and to take preliminary steps toward making it less ideological and more precise.

When quantum algorithms first appeared, it was argued that their physical realization is similar to analog computation and hence unfeasible because of susceptibility to noise (Landauer 1995; Unruh 1995). Quantum error correction (QEC) was conceived in order to meet this challenge, claiming that quantum computers are neither digital nor analog. Rather, several error-correction codes were developed that demonstrated that quantum computers are accurate continuous devices whose possible input, output, and internal states are dense countable sets (Shor 1995; Steane 1996). The next obvious worry was the cost of error correction: since this process is itself a form of computation, the expectation was that it would introduce additional noise to the computer. Fault-tolerant quantum computation was meant to address this concern. Several theorems have been proven (Aharonov and Ben-Or 1996; Knill and Laflamme 1996; Knill, Laflamme, and Zurek 1996, 1998) that rely on a set of requirements under which arbitrarily long computation can be fault-tolerantly executed on a quantum network if the error probability (or error rate) is below a certain threshold. Current estimations of the threshold point at an error probability of the order of 10^{-4} per computational step, which is still too low for existing technology. Given these results, the crucial open question is the following: Would improvement in hardware lead to the realization of large-scale, fault-tolerant, and computationally superior quantum computation, or is there still a fundamental obstacle that renders the entire project a wild goose chase?

Optimists within the quantum information community regard the challenge as purely technological, the answer to which is contingent upon the design of quantum gates whose accuracy is below the required error

https://doi.org/10.1086/648600 Published online by Cambridge University Press

threshold. Yet from a computer science perspective it is quite natural to approach the question with healthy pessimism, according to which, somewhere along the way, in the scaling of the computer from several qubits (the basic information-carrying units of the quantum computer) to a large collection thereof, the computer loses its "quantumness" and with it its putative computational superiority.[2] On this view there are in principle (as opposed to merely technological) reasons why there are no such computers, and no matter what gadgets are used, large-scale quantum computers will never be more efficient than their classical counterparts.

According to an influential view within the quantum information community (Aaronson 2004), this pessimism is tantamount to rejecting the universal applicability of quantum mechanics. But is this influential view correct? Must quantum computer skeptics reject the universality of quantum mechanics? Questioning one of the premises quantum computer optimists rely on, namely, the famous threshold theorems, this article argues they need not do so.

The key point is a striking yet rarely mentioned analogy that exists between the current "active" approach to FTQEC and the radical "interventionist" solution to the puzzle of thermodynamic arrow in time in the foundations of SM. In both approaches, external noise is seen as the culprit behind the irreversible behavior of physical systems governed by reversible dynamics. On this view, since all macroscopic systems are practically "open" and susceptible to noise, reversible behavior on macroscopic scales is unobservable unless the external noise is suppressed. In active FTQEC, noise suppression is achieved, among other things, by constantly "cooling" the computer, adding ancillas to the (open) quantum system. Remarkably, both approaches make a similar methodological choice in the way they treat noise in their models. This choice, or so I shall argue, is what makes the famous threshold theorems physically suspect.

The plan is the following. Section 2 presents the basic ingredients of active FTQEC, reconstructing it as an attempt to shield the computer from noise. Section 3 examines recent doubts with respect to the internal consistency of the assumptions behind active FTQEC. Section 4 reconstructs the debate on the feasibility of active FTQEC in a different context, taking its cue from an analogous debate on the thermodynamic arrow in time that exists within the foundations of SM. Based on this analogy and on the results from Section 3, Section 5 suggests that the active attempt

2. Of course, the fact that, so far, there is no agreed upon characterization of the physical resources behind the putative computational superiority of quantum algorithms over their classical counterparts, apart from, maybe, crude features such as 'entanglement', 'interference', or 'Hilbert space structure', to name a few, doesn't add to the clarity of this debate.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

to shield the computer from noise is the wrong way to think about FTQEC.

2. Optimism. In 1994, quantum computing, until then a speculative research area, received a dramatic boost. Built on an earlier discovery (Simon 1994) and on an ingenious number-theoretic argument, a quantum algorithm was presented (Shor 1994) that could factor integers into primes in polynomial number of computational steps—a subexponential speedup with respect to any known classical algorithm to date.[3] The discovery drew significant attention (such an algorithm, if implemented on a large-scale computer, could have dramatic consequences on current cryptography protocols) but was initially met with skepticism that portrayed the theoretician's dream as the experimentalist's nightmare (Haroche and Raimond 1996). Roughly speaking, this skepticism was based on the indisputable fact that large-scale quantum systems are noisy and on the observation that noise, or the accumulation of errors, would impede any attempt to physically realize a (computationally superior) large-scale quantum computer. In particular, Unruh (1995) was among the first to point out that the effects of the inevitable coupling of the putative quantum computer to external degrees of freedom put stringent constraints on the computation time, as the computer would lose its coherence (and also its "quantumness") in time scales that are determined by the strength of the coupling and the state of the environment.

One may wonder why noise was seen as such a threat. After all, there is a highly successful theory of classical error correction dating back to the early days of computer science that demonstrates how to protect computation against classical errors (von Neumann 1956).

The rub, of course, is that quantum mechanics is different than classical mechanics. In particular, classical errors are discrete by their nature, boiling down to bit flips or erasures. But a quantum state, represented, say, as a point on the Bloch sphere, is a priori continuous, and hence also the error is continuous. In a similar vein, quantum gates are continuous by their nature, never implemented exactly, but only up to a certain precision. Capitalizing on these differences, Landauer (1995) pointed out that small errors can accumulate over time and add up to large, uncorrectable errors. On this view, a quantum computation is nothing but analog computation, much more prone to noise, and thus actually inferior to digital classical computation.

Moreover, in order to protect against errors, in classical error correction one encodes the information in a redundant way, applying a majority

3. The best 'sieve' algorithms for FACTORING are subexponential. See, e.g., Lenstra and Lenstra 1990.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

check on the "dressed" bits that drastically reduces the probability of error.[4] But the unitarity of quantum mechanics, manifest in the quantum no-cloning theorem (Dieks 1982; Wootters and Zurek 1982) forbids one from copying an unknown quantum state without altering it and seems to prevent such a redundancy. Finally, classical error correction is based on the fact that one can acquire information on the type of error that had occurred (the "syndrome") without destroying the information encoded, but quantum mechanics seems to prevent one from doing so, since measurements collapse the quantum state and may destroy the information that was previously encoded in it. It thus came out as a complete surprise that shortly after these skeptical attacks there emerged a beautiful mathematical project (Shor 1995; Steane 1996) that demonstrated the possibility of quantum error correction.

2.1. The Birth of Quantum Error Correction. The gist of QEC lies in the quantum generalization of the classical repetition code. Suppose we would like to protect a qubit against a bit flip.[5] The protocol starts by encoding it into a larger space without copying or measuring it, that is, a pair of CNOT gates and a pair of ancilla qubits are used to implement the transformation:[6]

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

4. The simplest classical error-correction code is the repetition code: each bit is replaced by three of its copies ($0 \rightarrow 000$; $1 \rightarrow 111$), and a majority check can still protect against one bit flip. The protocol fails when two or more bits are flipped. However, if the error rate for one bit flip ε is assumed to be independent on each bit, then after the encoding the probability that we cannot correct a bit flip becomes $3\varepsilon^2(1-\varepsilon)+\varepsilon^3$ (there are three possible ways to have two bit flips and one way to have three bit flips), and so for $\varepsilon < 1/2$ we gain by encoding. See Kempe 2005, 6, and MacWilliams and Sloane 1977, 18–23.

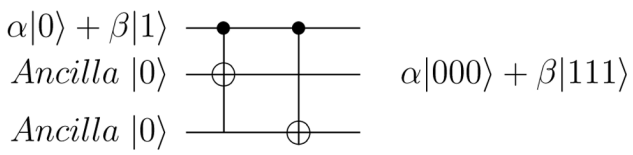
5. What follows in this subsection relies on an extremely clear exposition of FTQEC by Kempe (2005). Apart from the original papers from the 1990s, the interested reader can find many introductions to QEC, in varying degrees of accessibility, in, e.g., Nielsen and Chuang 2000, 425–499; Mermin 2007, 99–135; Gaitan 2008.

6. The CNOT gate is a two-qubit operator where the first qubit is the control and the second qubit is the target. The action of the CNOT gate is defined by the transformations

$$|00\rangle \rightarrow |00\rangle; |01\rangle \rightarrow |01\rangle,$$
$$|10\rangle \rightarrow |11\rangle; |11\rangle \rightarrow |10\rangle,$$

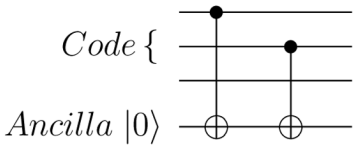
where $|00\rangle \equiv |0\rangle|0\rangle; |01\rangle \equiv |0\rangle|1\rangle$, etc. The first line of the transformation signifies that when the control qubit is in the "0"-state, the target qubit does not change after the action of the CNOT gate. The second line means that if the control qubit is in the "1"-state, the target qubit changes value after the action of the CNOT gate.

https://doi.org/10.1086/648600 Published online by Cambridge University Press



(2)

If a bit flip happens now, and the superposition in equation (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. For instance, we can measure the parity of the first two qubits with the following circuit (adding another ancilla qubit to the code and employing two additional CNOT gates):



(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.

Note that the parity measurement does not destroy the superposition. If the first qubit is flipped, then both $|100\rangle$ and $|011\rangle$ have the same parity (i.e., 1) on the first two qubits. If no qubit is flipped, the code word is still in the state of equation (1). This parity will be 0 for both $|000\rangle$ and $|111\rangle$.

Of course, a bit flip is but one of a continuum of possible quantum errors that a qubit can suffer. Indeed, in general it will undergo some unitary transformation in the composed system of qubit and environment that would entangle both. But it can be shown that any such unitary transformation that the composed system may undergo can be expressed as a linear combination of four basic errors known as the Pauli group, namely, bit flip, phase flip, a combination thereof, and the identity. In other words, these four error types span the space of unitary matrices that can affect the qubit.[7] By performing a measurement, we collapse the combined state on one of the four 'error subspaces' and hence disentangle the error from the information stored in the qubit without destroying it.

7. If we trace out the environment (average over its degrees of freedom), the resulting operator can be expanded in terms of the Pauli group, and we can attach a probability to each Pauli group element.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

This way, even though the quantum error is continuous, it will become discrete in the process of QEC.[8]

A more complicated encoding exists for phase-flip errors that uses nine qubits and can also correct a bit-flip error and a combination of both.[9] As in the case of classical error correction, the redundancy allows one to improve on the error probability for a single-qubit ε, since—with the discretization resulting from the error-recovery measurement—the state will be projected onto a state where no error has occurred with probability $1 - 9\varepsilon$, 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 $\leq 36\varepsilon^2$), the error is irrecoverable. Thus Shor's 9-qubit QEC is advantageous whenever $\varepsilon \leq 1/36$ (Kempe 2005, 8).

Soon after Shor's code appeared, several alternative QECs were suggested (e.g., Steane 1996; Gottesman 1997). It turned out that the smallest QEC that corrects a single error has five qubits and that this is optimal (Laflamme et al. 1996).

Before we go on to present the active fault-tolerant version of QEC, an important remark is in order. The crucial assumption employed in the error models that all these codes share is that the error processes affecting different qubits are independent of each other. This independence assumption is also present in the treatment of errors arising from imperfect gates or from state preparation and measurements. Moreover, the first noise models that were used in the proofs of the threshold theorems (see below) introduced another independence assumption, namely, an as-

8. In the above example, the parity measurement disentangles the code qubits from the environment and acquires information about the error. The three parities (for each qubit pair of the code word) give complete information about the location of the bit-flip error. They constitute what is called the *error syndrome measurement*. The syndrome measurement does not acquire any information about the encoded superposition, hence does not destroy it. Depending on the outcome of the syndrome measurement, we can correct the error by applying a bit flip to the appropriate qubit.

9. Shor's idea was to encode a qubit using nine qubits in the following way:

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

$$|1\rangle_{enc} = 1/\sqrt{2^3}(|000\rangle - |111\rangle)(|000\rangle - |111\rangle)(|000\rangle - |111\rangle).$$

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. 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

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

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).

https://doi.org/10.1086/648600 Published online by Cambridge University Press

sumption about the Markovian character of the qubit-environment interaction. This means that the environment maintains no memory of the errors, which are thus uncorrelated in time as well as in qubit location. As we shall see, this assumption is at the heart of the debate on the feasibility of active FTQEC.

2.2. The Miracle of Active Fault-Tolerant QEC. Admittedly, error recovery is a type of computation, hence it is not a flawless process; in the course of error correction, the application of the code itself (e.g., the parity measurements above) might suffer from imprecision. Such errors, if present, would propagate with the recovery process, leading to its inevitable corruption. Computation, moreover, requires not only storing information but also manipulating it. One could decode the (error-corrected) qubits, apply quantum gates on them, and reencode them, but this means exposing the qubits to additional noise and decoherence. Consequently, one must perform quantum gates directly on the encoded, or "dressed," qubits.

By following specific guidelines for preventing error propagation in the error-recovery process (i.e., verifying the ancillas, encoding the qubit, detecting the error syndrome, and recovering the error), one can demonstrate that fault-tolerant storage can be achieved (Shor 1996; Preskill 1998). With all the precautions, recovery will only fail if two independent errors occur in this entire procedure. The probability that this happens is still $c\varepsilon^2$ for some constant c that depends on the code used.[10] Additional work has shown how to apply the gates directly to the encoded data, without introducing errors uncontrollably, and several variants of fault-tolerant universal quantum computation have been developed for different QEC codes (Gottesman 1997).

Thus, in the few years that followed the introduction of Shor's algorithm, we have seen the emergence of remarkable and beautiful mathematical results that demonstrate how to encode quantum data redundantly, how to perform fault-tolerant data recovery, and how to compute fault-tolerantly on encoded states. This seems too good to be true: quantum codes exist that can correct up to t errors, where t can be as large as we wish, and on which we can compute fault-tolerantly. This means that if our error rate is ε, then computation will only fail with probability of order ε^{t+1} for a t of our choice.

From the skeptical point of view that we are interested in here, however, the issue at stake is the actual cost of QEC. If this cost scales too fast in terms of computational resources, then one has achieved nothing by em-

10. This constant can be quite large because there are now many more gates and steps involved in maintaining the process reliable.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

ploying it, as these resources would offset the putative superiority of the ideal quantum computation (without the noise). Put differently, 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. Yet a quick calculation shows that 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.[11]

In a remarkable feat of mathematical ingenuity, this obstacle was overcome by using concatenated codes that involve recursively reencoding already encoded bits (Knill and Laflamme 1996).[12] More precisely, in a first-level encoding one encodes each qubit with an appropriate code. Then, for each of the code words one encodes each of the qubits again using the same code. This process, which scales only polylogarithmically with the length of the original (ideal) algorithm,[13] reduces the effective error rate at each level, with the final accuracy being dependent on how many levels of the hierarchy are used.[14] Moreover, it was shown that 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 any quantum computation with arbitrary accuracy.

Shifting the problem to the technological realm (i.e., the design of quantum gates below a certain error threshold, which was estimated to be of the order $10^{-4}-10^{-6}$ errors per computational cycle), optimists could now declare:

> Therefore, noise, if it is below a certain level, is not an obstacle to unlimited resilient quantum computation. (Knill et al. 1998, 342)

11. The number of steps required for recovery scales as a power of t, t^a with exponent $a>1$. That means that the probability to have $t+1$ errors before a recovery step is completed, scales as $(t^a\varepsilon)^{t+1}$. This expression is minimized when $t=c\varepsilon^{-(1/a)}$ for some constant c and its value is at least $p=\exp(-ca\varepsilon^{-(1/a)})$. This means that our probability to fail per error correction cycle is at least p. If we have N such cycles, our total failure probability is $Np=\exp(-ca\log N\varepsilon^{-(1/a)})$. For $p\ll 1$, ε must scale as $(1/\log N)^a$. See also Kempe 2005, 13.

12. This idea has, again, classical roots, e.g., MacWilliams and Sloane 1977, 307–316.

13. If the number of gates in the original algorithm is N, the fault-tolerant version thereof has $N\,\mathrm{poly}(\log N)$ gates.

14. For example, the two-level concatenation reduces the error probability per gate p from cp^2 to $c(cp^2)^2=c^3p^4$ for some constant c that depends on the code, which 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 $(cp)^{2^k}/c$, which means the error rate decreases faster than the size of the circuit grows. See, e.g., Kaye, Laflamme, and Mosca 2007, 237–239.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

This optimism is still present today in many expositions of FTQEC:

> 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, 240)

> The [threshold] theorem made it clear that no physical law stands in the way of building a quantum computer. (Gaitan 2008, xvii)

Pessimists, however, were not convinced.

3. Pessimism. The threshold theorems were an important landmark in the creation of an industry, as they shifted the debate on the feasibility of large-scale, fault-tolerant, and computationally superior quantum computation to the technological domain.[15] But the assumptions that were crucial to the above remarkable results came under scrutiny only several years later.[16] We have already encountered one of these assumptions:

(A) Error correlations decay exponentially in time and space.

This assumption basically means that the environment, whose state is traced over in the computational process, maintains no memory of its states over time and, moreover, that the error model that is assumed in the estimation of the threshold involves spatially uncorrelated errors. Both these features, mind you, are in stark contrast to the correlations that characterize the system whose qubits are supposed to be entangled.

Among the additional assumptions that are necessary for the validity of the threshold theorems, two are crucial to the analysis that follows:[17]

15. As an aside, it is my opinion that the sociological issues surrounding the rise of this industry would one day be a worthy subject for historians and sociologists of science, e.g., the fact that most of the purely mathematical results of the threshold theorems that were taken for granted, at least at the early stages, and later disseminated as "mantras," were produced by mathematicians and computer scientists, and apart from having little to do with actual physics, were impenetrable to most working physicists.

16. The following section relies on an analysis based on the theory of open quantum systems suggested by Alicky, Lidar, and Zanardi (2006).

17. There are other necessary assumptions, e.g., the ability to perform parallel operations, or a constant error rate, that shall not be discussed here.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

(B) Gates can be executed in time τ_g such that $\tau_g\omega = O(\pi)$, where ω is the Bohr or the Rabi frequency.
(C) A constant supply of 'fresh', nearly pure, ancilla qubits is available.

While not stated explicitly in the FTQEC literature, assumption (B) follows from the definition of a quantum gate, which is basically a unitary transformation $U = e^{iA}$, where $A = \tau_g H$ and H is the Hamiltonian generating the gate. When τ_g is scaled up, H (and hence its eigenvalues) must be scaled down, and vice versa.[18] Assumption (C), however, is stated explicitly (Aharonov et al. 1996) and represents the ability to continuously dump the entropy produced during the noisy computational process.

The problem, however, is that there appears to be an internal inconsistency between these three assumptions: when one inquires into the derivation of the Markovian master equation that characterizes the dynamical evolution of an open quantum system (which is presumably the case at hand according to assumption [A]), one finds that the two coupling limits between the quantum computer and its environment (the singular limit and the weak limit; see Appendix) impose constraints that are not satisfied by the other two assumptions. In particular,

- (A) and (B) are incompatible with the weak coupling limit and so require the singular coupling limit, which means that the reservoir (the source for the ancillas) must possess a high temperature, which then contradicts (C).
- (A) and (C) are incompatible with the singular coupling limit and so require the weak coupling limit, which means that the gate velocity must be slow, which then contradicts (B).

We shall now look more closely at this inconsistency claim.

Active FTQEC, and in particular the threshold theorems proven in 1996–97, started from a phenomenological point of view, assuming a very simple, "well-behaved," error model: "We assume that a gate's error consists of random, independent applications of products of Pauli operators with probabilities determined by the gate" (Knill 2005, 40). This approach was natural given the difficulty in obtaining thresholds for models that

18. The distinction between the Bohr or the Rabi frequencies depends, respectively, on the application of constant vs. periodic fields in controlling the gate. For our purposes, however, in both cases the crucial insight is that the gate "velocity" is directly related to the energy gap between the ground and the excited states of the Hamiltonian that generates the gate: fast gates require large energy gaps, and vice versa. See also Margolus and Levitin 1998 for a lower bound on the energy required to transform a state $|\psi\rangle$ to its orthogonal using a constant Hamiltonian, and Gea Banacloche 2002 for a lower bound on the amount of energy needed to carry out an elementary logical operation in a quantum computer.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

are not phenomenological but that start instead from a purely microscopic, Hamiltonian description. Such a description, however, cannot be ignored, and, as I shall argue here, may also lead to some interesting philosophical insights.

Assumption (A) arises from the upper bound on the probability p_{err} of a faulty computational path with k errors, required by the threshold theorems (Aharonov et al. 1996, Section 2.10):

$$p_{\text{err}} \leq c\varepsilon^k(1-\varepsilon)^{v-k}, \tag{4}$$

where ε is the probability of a single error, c is a certain constant independent of ε, and v is the number of error locations in the circuit. This bound implies that the k-qubit errors should scale as $\sim\varepsilon^k$, and this can be satisfied only for reservoir correlation functions whose temporal decay is exponential.[19]

From a dynamical perspective this situation can be investigated with the derivation of the quantum Markovian master equation (MME; see the appendix). When such an investigation is carried out, the following points emerge:

- There are two types of fully rigorous derivations of quantum MME, known as the singular coupling limit (SCL) and the weak coupling limit (WCL).
- Both derivations must satisfy a thermodynamic constraint, namely, that the reservoir is in a state of thermal equilibrium.[20]
- Within the SCL, this condition allows the reservoir's correlation function to be approximated by a delta function (a necessary condition for obtaining the Markovian approximation), which is the case only in the limit of infinite temperature.
- Thus, within the SCL, a Markovian approximation that satisfies the thermodynamic constraint can be justified only in the limit $T_R \to \infty$, where T_R is the reservoir's temperature.
- Physically, this means that the 'zero-memory' condition that is encapsulated in assumption (A) holds within the SCL only when the reservoir is much hotter than the system on the same energy scale set by the system + ancillas.

The final point, while consistent with assumption (B) (as it allows arbitrary gate velocities), appears to contradict assumption (C): where does one get 'fresh, almost pure' ancillas to dump entropy in if the reservoir's

19. The localization of errors in time translates into localizations of errors in space due to the finite speed of error propagation.

20. This constraint is known as the Kubo-Martin-Schwinger (KMS) condition; see Kastler 1978.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

temperature is much higher than the temperature of the system? Inversely, if one requires pure ancillas, then by coupling them to the system one must abandon the Markovian noise model in the environment.[21]

Now one may argue that the SCL, while mathematically rigorous, is nevertheless highly unphysical due to the singular character of the interaction and the infinite temperature condition. Yet the more realistic domain of WCL appears to be as unfavorable to FTQEC as the SCL:

- Within the WCL (where the reservoir's temperature is finite), one can achieve the Markovian condition in the reservoir's correlations function only after coarse graining over very long time scales.[22]

Expressed in terms of the system's gates frequency, this condition, while consistent with assumption (C), violates assumption (B), as it only allows slow, adiabatic gates.[23]

The Markovian approximation within WCL is thus a long time limit (at least compared to the system's gate frequency), and so one cannot expect it to hold in the time scales that are stated in the assumptions behind active FTQEC. If one wants to fault-tolerantly implement gates with a finite speed, one must, again, abandon the simple uncorrelated noise model and consider non-Markovian noise.[24]

It is noteworthy that the above considerations are unavoidable if one accepts thermodynamics, as they follow from a derivation of the MME that satisfies the KMS condition (see the appendix), which is basically a restatement of the minus-first law (Brown and Uffink 2001), namely, that a system driven out of equilibrium relaxes toward a new equilibrium state

21. Recall that the role of fresh ancillas is to remove the entropy that is generated from the system-bath interaction. If the ancillas are already hot, this means they are initially entangled with the system, and this introduces more errors into the computational process than QEC could handle.

22. Davies 1974; this famous derivation of the quantum MME in the WCL is perhaps the only one consistent both mathematically and physically.

23. Further analysis is possible according to the type of the driving Hamiltonian (constant or periodic) and the Bohr or the Rabi frequencies, respectively. This, however, doesn't change the requirement for adiabatic gates.

24. Relaxing the requirement for fast gates without abandoning the Markovian condition will only generate more inconsistency: as we have seen, within the WCL the more adiabatic the evolution, the smaller the probability for correlated noise per gate, but in active FTQEC, this probability is also inversely related to the input size. This means that if we work with adiabatic gates in active FTQEC with a Markovian noise model, we violate one of the thresholds conditions which states that the gate velocity and the input size are independent.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

set by the external constraints.[25] Now one can, of course, doubt the validity of thermodynamics (and in particular the applicability of the minus-first law) in certain nongeneric situations and argue that a quantum computer is such a case. Here we note that such a claim was not made within the FTQEC community.[26] Rather, instead of abandoning active FTQEC, a shift was made toward proving threshold theorems for non-Markovian noise (Terhal and Burkard 2005; Aharonov et al. 2006). Prior to considering this shift, let us take a short historical detour.[27]

4. Those Who Cannot Remember the Past. In this section I would like to suggest that the original threshold theorems that are based on uncorrelated noise supply yet another example of Santayana's Aphorism on Repetitive Consequences,[28] and could have been deemed of little physical significance much earlier. Support for this claim comes from a historical analogy that exists between active FTQEC and a certain radical interpretation of a controversial school in the foundations of SM. Combined with the above results, this analogy raises doubts whether active FTQEC is the correct way to approach the realization of a large-scale, fault-tolerant, and computationally superior quantum computation.

4.1. Two Problems in the Foundations of SM. That many things are easier to do than to reverse has been known to mankind from time immemorial. In contrast, the fundamental laws of physics are believed to indifferent to the directionality in time, allowing in principle such a reversal. One of the great challenges to mathematical physics in the twenty-first century is thus to construct a rigorous mathematical derivation of equations that describe irreversible behavior, for example, Fourier's law of heat conductivity, from a classical or quantum model with a Hamiltonian microscopic dynamics (Bonetto, Lebowitz, and Rey-Bellet 2000; Alicki and Fannes 2001). In this respect the fundamental question that the founding fathers of the kinetic theory struggled with in the second

25. For this reason the common argument, according to which a quantum computer—and in fact any computer—is not in an equilibrium state, so equilibrium thermodynamics doesn't apply, is completely beside the point. That this is a nonstarter is evident from the fact that the equilibrium condition is imposed only on the heat bath, and not on the system. Moreover, using many baths with different respective temperatures is of no avail, as this doesn't resolve the inconsistency.

26. I defer the discussion on this line of thought (and on the reason why it is not an option for the FTQEC community) to future work.

27. If taken 10 years earlier, this detour might have saved the quantum information industry some time.

28. "Those who cannot remember the past are condemned to repeat it" (Santayana 1905, 284).

https://doi.org/10.1086/648600 Published online by Cambridge University Press

half of the nineteenth century remains open: Can a reversible and deterministic dynamics fully explain the behavior of macroscopic matter? The answer is not so obvious, as by now we have many examples where it is possible to derive rigorously the transport laws by adding a non-Hamiltonian component—either stochastic or deterministic—to the microscopic Hamiltonian dynamics (for the former, see, e.g., Davies 1978; for the latter, see, e.g., Gallavotti and Cohen 1996), while the only Hamiltonian models for which such a rigorous proof exists are highly idealized and involve noninteracting particles (Lebowitz 1978; Lanford 1981).

Another famous problem in the foundations of SM is to account for the statistical assumptions introduced into the underlying dynamics that presumably govern thermodynamic phenomena[29] and in particular for the assumption of equiprobability (if all states consistent with some global equilibrium constraint are equiprobable—i.e., if the probability distribution in phase space is uniform over the constraint surface—then by averaging over them one reproduces the thermodynamic relations).[30] Roughly there are three main approaches to this problem. One can regard this assumption as a priori true, justifying this belief with the principle of indifference. There are two obvious difficulties in this position: (1) What is the meaning of an average when one deals with an individual system? (2) How can a subjective notion of probability affect objective notions such as entropy? One can follow Gibbs and consider an ensemble of systems, but while this approach is formally consistent (and also widely used), it does not resolve the first difficulty; there is usually only one actual system. Alternatively, one may follow Boltzmann's ergodic hypothesis and try to dynamically justify equiprobability, but ergodicity holds only in infinite time limits, and besides, many (thermodynamic) systems are not ergodic.[31]

An important but controversial alternative to these approaches within the foundations of SM is the interventionist school, also known as the open system approach. In its premise lies an undeniable fact: all physical system are "open," in the sense that they interact with their environment.

29. That this problem is independent of the former sometimes goes unnoticed. See Price 2003 for such a mistake, and Hagar 2005 for a criticism.

30. As an anonymous referee remarks, there are further complications here, e.g., a distribution that is uniform with respect to one set of variables need not remain so under reparameterization, and some expalnation, presumably dynamical, is required for favoring a distribution that is uniform under canonical coordinates, rather than some other distribution.

31. Recent results in quantum thermodynamics (e.g., Tasaki 1998; Gemmer, Michel, and Mahler 2004; Goldstein et al. 2006; Popescu, Short, and Winter 2006) suggest to replace the assumption of equiprobability with a simple and natural assumption about quantum entanglement.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

But while this fact is acknowledged as trivial and benign by all the competing approaches to the foundations of SM, within the interventionist school it becomes the key to the solution of the two foundational problems mentioned here.[32]

4.2. The Roots of the Open System Approach. Interventionist ideas appear already in the discussions on the plausibility of the probabilistic assumptions behind Maxwell's derivations of the velocity distribution and in Boltzmann's formulation of ergodicity. For example, Maxwell mentions "interactions with the surroundings" as a possible justification for his equipartition theorem (quoted in Brush 1976, 366–367), and Boltzmann ([1895] 1964, 60) points out that the special (highly ordered) initial conditions that lead to abnormal thermodynamic behavior, that is, to entropy decrease, "could be destroyed at any time by an arbitrarily small change in the form of the container."

While Maxwell and Boltzmann were ambivalent with respect to the character of these external perturbations, Boltzmann's advocates in the debate on the validity of his *H*-Theorem that took place on the pages of *Nature* between 1894 and 1895 explicitly claimed that random external perturbations lead to the desired randomization of molecular motion—a randomization that was at the heart of Boltzmann's derivation of Maxwell's velocity distribution from Hamiltonian dynamics. Here is Burbury, an English barrister who took on himself the task of defending Boltzmann's views: "Any actual material system receives disturbances from without, the effect of which coming at haphazard, is to produce the very distribution of coordinates which is required to make *H* diminish, so there is a general tendency for *H* to diminish, although it may conceivably increase in particular cases, just as in matters political change for the better is possible but the tendency is for all change to be from bad to worse" (1895, 320).

Several years later Borel explained why external interventions should be instrumental in the construction of mechanical models of thermodynamic phenomena: "The representation of gaseous matter composed of molecules with position and velocities which are rigorously determined at a given instant is therefore a pure abstract fiction; as soon as one supposes the indeterminacy of the external forces, the effect of collisions will very rapidly disperse the trajectory bundles which are supposed to be infinitely narrow, and the problem of the subsequent movement of the molecules becomes, within few seconds, very indeterminate, in the sense that an enormously large number of different possibilities are a priori

32. For an interesting defense of a benign version of the open system approach, see, e.g., Shenker 2000.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

equally probable" (Borel 1914). According to Borel even the gravitational effects resulting from shifting a small piece of rock with a mass of one gram as distant as Sirius by a few centimeters would completely change the microscopic state of a gas in a vessel here on Earth by a factor of 10^{-100} within a fraction of a second after the retarded field of force has arrived.[33]

Note that Boltzmann and Borel both limit the effect of the external intervention to the instability of those initial microstates that lead to abnormal thermodynamic behavior. However, it seems unfair to bring external perturbations as deus ex machina to save the foundations of SM and then neglect them in the subsequent dynamics (Earman 1986, 231). Interventionist models—I shall call these "radical" henceforth—that have appeared in the 1950s have subsequently modified the dynamics of thermodynamic systems to include random external perturbations,[34] arguing that these perturbations are absolutely necessary for the construction of realistic mechanical models for thermalization: "Statistical mechanics is not the mechanics of large, complicated systems; rather it is the mechanics of limited, not completely isolated systems" (Blatt 1959, 749).[35]

4.3. There Is Always a Little Noise. Those interventionists who interpret the open system approach "radically" reveal two fundamental features in their models. First, viewing their approach as a solution not only to the problem of equiprobability but also to the problem of irreversibility, they regard external noise as necessary for the validity of thermodynamics and, in particular, for the process of thermalization. If this noise is removed, no such process can take place.[36] Second, within this radical ap-

33. Such perturbations have dramatic consequences for a system whose dynamics are unstable. Michael Berry (1978, 95–96) has calculated that an electron in the limits of the observable universe will disturb the motion of two colliding oxygen molecules here on Earth to the extent that predictability of the motion would be lost after 56 collisions. Note that we should not go as far as the end of the universe. Even a medium size billiard player situated 1 meter far from a billiard table would do the same for two elastic billiard balls after nine collisions.

34. Famous milestones after Borel's calculation are the Spin Echo experiments (Hahn 1950) and the subsequent models they initiated (e.g., Lebowitz and Bergmann 1955; Blatt 1959). A revival of this approach has been recently championed by Earman (1986), Earman and Redei (1996), and Ridderboss and Redhead (1998).

35. In his model Blatt nicely demonstrates how just by taking into account the walls of a container in which a gas approaches equilibrium one can achieve realistic time scales for this process.

36. Radical interventionist models assign the environment an active part; as such they go beyond the usual role the environment plays in thermodynamics. For an analysis, see Uffink 2001.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

proach there is an unexplained dichotomy between the system and its environment; while the former is treated as deterministic, the latter is allowed to include random elements.

Agreed, some interventionists may try to come up with an explanation to this dichotomy, arguing that, for example, compared with the system, the environment is large and complex, and so its influence on the system is effectively random, or that the correlation established between the system and the environment during the interaction play little role in the subsequent evolution; but the crucial point here is that these explanations matter little to the conceptual problems posed in the beginning of this section (namely, the problem of irreversibility and the problem of equiprobability). First, if the underlying dynamics (of the open system and its environment) are time-reversible-invariant, then the perturbations (or noise) can be reversed. Second, by describing these perturbations as random, one presupposes the very randomness one would like to underpin in the first place within one's system (Hagar 2005, 476).

The open system approach was rediscovered by the decoherence school in quantum theory that lately has become the new orthodoxy.[37] There are direct historical links between interventionist models and the decoherence school,[38] but more important here is the conceptual link. In particular, the two features identified above, namely, the role of noise in thermalization processes and the double standard in the characterizations of the system and the environment, play a crucial part in the theory that was developed to suppress decoherence. This theory, as we recall, is active FTQEC.

What makes active FTQEC special—and what is emphasized in the literature on the threshold theorems—is the remarkable feat of rendering QEC, the process of overcoming the noise (by discretizing it, by disentangling it from the data through its entangling to ancillas, by dumping the ancillas, and by constantly supplying fresh ones; see Section 2), an inexpensive process from a computationally complexity perspective. In other words, whatever cost this process may incur, the threshold theorems guarantee us that below a certain error rate, when "balancing the books," we will not transcend the putative complexity barrier between the quantum and the classical: the cost of QEC will not increase exponentially with the input size, and the quantum computer would maintain its alleged computational superiority.[39]

37. The term was coined by Bub (1997).

38. Joos et al. (1996) cite Borel (1914) in their introduction.

39. Or, put differently, the protocol will be fast enough to overcome the error rate so as to avoid errors accumulation that will bring us back to a classical computational power.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

We can now make the analogy more explicit: think of the quantum computer as an open system out of thermodynamic equilibrium, and on the noise as arising from the computer's interaction with a heat bath. Active QEC in this context can be seen as the attempt to "cool down" the open system, thus preventing its thermalization, and the threshold theorems are the promise that given a certain noise level, this "prevention" can be done for an arbitrarily long time with only moderate overhead, that is, 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.

Now there are at least two reasons to be suspicious of this marvelous result, a result that, among other things, renders the concept of a "one-way" function (a function that is easy to compute but infeasible to reverse) obsolete.[40] The first is the latent double standard with which the system and the environment are treated here: while the interaction between the qubits + ancillas is entangling (and so nonlocal correlations dynamically evolve), in the interaction with the environment (or the noise model) no nonlocal correlations are allowed to evolve. In particular, with every computational step the environment acts as if it has "seen" the system for the first time. Note that classical error correction incorporates no such double standard since the dynamics that governs the bits, as well as the noise are both assumed to be yield no entanglement from the outset. Given this double standard, the original threshold theorems seem now less miraculous: if one is allowed to cheat just once in quantum mechanics, one can indeed do miracles.[41]

The second reason to be suspicious is that active FTQEC operates under the tacit assumption that external noise is responsible for the thermalization of the quantum computer: "Quantum error-correction may be thought of as a type of a refrigeration process, capable of keeping a quantum system at a constant entropy, despite the influence of noise processes which tend to change the entropy of the system" (Nielsen and

40. "One way" functions form one of the most productive discoveries in computer science. Many results in modern cryptography—especially the celebrated RSA cryptographic protocol (Rivest, Shamir, and Adleman 1978)—depend on the infeasibility of recovering x from $f(x)$. In particular, while any arbitrary variable x contains as much (Kolmogorov algorithmic) information as the function $f(x)$, proving such an equivalence requires the ability to perform an exhaustive search; see, e.g., Kolmogorov 1965 and Levin 2003.

41. Paraphrasing Rovelli's (1994, 2790) famous attack on the notion of "protective measurement." Analogously, if one inserts just one ideal (i.e., nonsensitive to thermal fluctuations) element deep in a complicated construction, one can easily use this construction to violate thermodynamics.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

Chuang 2000, 569). Again, such a claim is open to the radical interpretation that it is only the noise that is responsible for thermalization, which means that if we eliminate the noise, no such process would take place.[42] But surely this cannot be true. If the composed system (i.e., the quantum computer coupled to the heat bath) were left to itself, it would eventually equilibrate regardless of external perturbations, wouldn't it? Well, thermodynamics tells us that all physical systems out of equilibrium do so (thus obeying the minus-first law), and statistical mechanics 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.

Note however, that if proponents of active FTQEC accepted this benign version of interventionism, then the motivation for looking at error correction as an active process would be somewhat blunted. In fact, if proponents of active FTQEC agreed to limit the consequence of external noise to that of perturbing the system form a state that exhibits abnormal thermodynamic behavior and putting it back on a state that exhibits normal thermodynamic behavior, then the question one should ask in the context of FTQEC is not how to eliminate the noise, but rather how to create those abnormal states to begin with.

These considerations lead one to rethink the conceptual framework of FTQEC. If the aim is to forestall thermalization of the quantum computer, maybe active FTQEC is the wrong way to go about it. After all, one of the lessons of the foundations of SM is that if one insists on keeping the dynamics time-reversal-invariant, the reason for thermalization cannot reside in the dynamics alone and must involve the system's state space. Thus, rather than trying to actively shield the computer from external noise, perhaps one should try to create those states that resist thermalization. Such states do exist; they are allowed by theory, and (as our experience shows) in rare occasions can even be prepared in the lab.[43] In future work I shall discuss this strategy and the skepticism associated with it. My goal here is just to draw attention to this perspective.[44] Let us now

42. I have found only one place where this contrapositive assumption has been explicitly made in print—not in a refereed publication, but in a blog entry from 2005; see http://dabacon.org/pontiff/?p=959.

43. See, e.g., the Spin Echo experiment of Hahn (1950).

44. This passive perspective bifurcated early on from active QEC (Barenco et al. 1996; Zanardi and Rasetti 1997). Whether or not this is how QEC is looked at now (e.g., Knill 2006) is an open question. It is noteworthy that the standard formalism for QEC, known as the stabilizer formalism (Gottesman 1997) provides a general framework for both active and passive QEC via its group theoretic form. But one should recall here that, (1) in spite of the formal resemblence, the two strategies are physically very

https://doi.org/10.1086/648600 Published online by Cambridge University Press

return to the current state of affairs in active FTQEC, and to the attempts to prove threshold theorems for non-Markovian noise.

4.4. Toward More Realistic Noise Models. Combined with the attack on the internal consistency of the assumptions behind the original threshold theorems, the historical analogy proposed here suggests that the error model introduced in these theorems was inadequate. Serving its purpose in yielding provable thresholds, it nevertheless failed to characterize correctly the system-environment interaction that active FTQEC purports to describe. It is noteworthy that already in his seminal paper on classical error correction, von Neumann admitted that the statistical independence of the error per bit is "totally unrealistic" (von Neumann 1956, 90). Indeed, the exponential decay of correlations of errors seems to be an idealization of both real internal fluctuations and real random external disturbances; every real process has some intrinsic built-in delay, which means the presence of memory, or non-Markovianity.

Non-Markovianity means that instead of thinking of noise as local, that is, as acting on the qubits independently of the structure of the evolution of the quantum computer (in which case this evolution only propagates the errors), the environment now "sees" the evolution of the computer and "learns" it. Since this evolution is nonlocal (recall that to protect the information from noise QEC must produce multiqubit entanglement), it will eventually give rise to nonlocal noise (Alicki et al. 2002, 8). Although the strength of this effect can be mitigated by lowering the velocity of the quantum gates, that is, by increasing the overall computational time (Alicki et al. 2002, 10–11),[45] 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 nonlocal is the noise.

What would be the effect of nonlocal noise on active FTQEC? This question is still open. First and foremost, the claim that "arbitrarily long" quantum computation can be done if the noise is below a certain threshold must be reconsidered. True, vacuum-induced decoherence is rather small in comparison to other sources of noise with which active FTQEC can deal, but the time scale in which the latter can now take place is certainly not "arbitrarily long," and, even more important, any statement about

different and (2) the threshold theorems discussed here were proven solely in the context of active FTQEC.

45. An illustrative analogy here is a swimmer immersed in a viscous fluid; the faster her strokes, the stronger is the fluid's resistance.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

its length requires a delicate analysis of different time scales (i.e., the time scale in which the memory is exponentially decaying vs. the time scale in which the vacuum memory prevails).

In particular, there are two crucial consequences for the threshold estimations. First, non-Markovian noise would constrain the threshold theorems to deal with amplitudes and not with probabilities in the calculations of the error rates, as the qubits and the environment would remain entangled for longer time scales. Since a probability is a square of an amplitude, this would have damaging effect on the error rate thresholds, as they would now become much lower than the previous (uncorrelated) case, making the entire project of active FTQEC physically unfeasible (Alicki 2004). One may argue that this estimation is highly pessimistic as it allows the bad fault paths to add together with a common phase and thus to interfere constructively. Most likely, or so the argument goes (Preskill 2007; Khoon and Preskill 2008, 6), distinct fault paths would have only weakly correlated phases, and if so, then the modulus of a sum of N fault paths should grow like $\sqrt{N}$ rather than linearly in N, leading to better threshold estimations. This more optimistic hypothesis remains to be proven, yet it is hard to see how one can justify such an intricate behavior of the noise phases on an a priori basis, especially when one requires the exact opposite behavior for the qubits (see Section 4.3).

Another problem that arises in threshold theorems for correlated (non-Markovian) noise (Terhal and Burkard 2005, Section 5) is that they explicitly rely on the norm of the interaction Hamiltonian $||H_I||$ (see Appendix). Now, apart from the fact that this quantity is not directly measurable in experiments,[46] the requirement for low error rate means, physically, that one assumes that the very-high-frequency component of the noise is particularly weak, a requirement that seems not to be physically well motivated, and in some decoherence models even implies that the system and the environment are practically decoupled, which renders the use of QEC obsolete (Alicki 2007; Hines and Stamp 2007; Khoon and Preskill 2008, 7).

Taking stock, while the jury is still out with respect to the feasibility of active FTQEC in the presence of non-Markovian noise, one thing is certain: 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,[47] while for other it may still

46. In fact, for otherwise reasonable noise models, the norm of the interaction Hamiltonian could be formally infinite (if, e.g., the qubits couple to unbounded bath operators).

47. For such an adversarial noise model, see Kalai 2005, 2006, 2008, 2009.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

be within reach for a certain amount of time. The irony is that the attempt to characterize actual quantum noise requires exponential resources (Emerson et al. 2007). It seems that we need a quantum computer to tell us whether a large-scale, fault-tolerant, and computationally superior quantum computer is possible.

5. Optimal Skepticism. Quantum information scientists would like to believe that what prevents them from building a large-scale, fault-tolerant, and computationally superior quantum computer is a technological barrier. Their optimism is based on the premise that since quantum correlations are stronger than classical correlations, computationally superior quantum computers are possible in principle, and so it is up to us to surpass the experimental obstacles involved in controlling the noisy evolution of a large-scale collection of entangled qubits. Consequently, quantum information scientists characterize pessimism about the feasibility of large-scale, fault-tolerant, and computationally superior quantum computers as ideological, the claim being that such a pessimism must go hand in hand with skepticism about the universal applicability of quantum theory itself (Aaronson 2004). Note, however, that if this were true, then the inability to create these fascinating machines would be judged rather harshly by the public.[48]

Here I have suggested that the puzzle of the feasibility of large-scale, fault-tolerant, and computationally superior quantum computers resembles the puzzle of the thermodynamic arrow in time: in both domains certain states or processes that are allowed by theory remain mostly unobserved, and in both domains there exists a debate on the extent to which external perturbations are responsible for this fact.

Given the above analysis of the threshold theorems and the analogy between active FTQEC and the radical interventionist school in the foundations of SM, it now appears that 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. As we have seen, the feasibility of FTQEC is dependent not only on technology, but also on the actual noise, and since the distinction between the system and the noise is completely arbitrary from a fundamental perspective, active FTQEC might be the wrong way to approach

48. A similarly severe constraint on skepticism with respect to the Large Hadron Collider is criticized by Dyson (2009). Dyson worries that portraying the failure of the LHC as the end of theoretical physics would be damaging to the latter in case the former does fail for some reason, contingent or fundamental.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

the project. In fact, active error correction seems to be a misnomer: the essence of the project should be passive rather than active; errors should be avoided and not corrected.

The final lesson is this. For the purpose of advancing the debate on the feasibility of large-scale, fault-tolerant, and computationally superior quantum computers, 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 contingent upon technological capabilities). Such a skepticism would be optimal as a counterexample to the received wisdom in the quantum information community. Here I have only argued that such a skepticism is possible, and have suggested that it should be developed in the context of passive FTQEC. Making this skepticism precise is, I believe, what pessimists about the feasibility of large-scale, fault-tolerant, and computationally superior quantum computers should aim for.

Appendix: The Dynamics of an Open Quantum System

What follows is based on standard introductions to the theory of Quantum Open Systems (Breuer and Petruccione 2003). In general terms an open quantum system $\mathcal{S}$ with a Hilbert space $\mathcal{H}_S$ is coupled to another quantum system $\mathcal{B}$, the environment, with Hilbert space $\mathcal{H}_B$ (if $\mathcal{B}$ has infinite degrees of freedom, it is sometimes called a 'reservoir', and if the reservoir is in thermal equilibrium, it is denoted a 'heat bath'). Thus $\mathcal{S}$ is a subsystem of the total system $\mathcal{S} + \mathcal{B}$ living in the tensor product space $\mathcal{H}_S \otimes \mathcal{H}_B$. If the interaction between $\mathcal{S}$ and $\mathcal{B}$ (with their respective Hamiltonians H_S and H_B) is governed by the Hamiltonian H_I, then the Hamiltonian of the total system can be written as

$$H = H_S \otimes I_B + I_S \otimes H_B + \alpha H_I, \tag{A1}$$

where I_S and I_B denote the identities in the Hilbert spaces of the system and the environment, respectively, and α is a coupling constant.

The open system $\mathcal{S}$ is singled out by the fact that all observables A of interest refer to this system, and are thus of the form $A \otimes I_B$, where A acts in $\mathcal{H}_S$. If we describe the state of the total system by some density matrix ρ, then the expectation value of the observable A is determined by

$$\langle A \rangle = tr_S\{A\rho_S\}, \tag{A2}$$

where

$$\rho_S = tr_B\rho, \tag{A3}$$

https://doi.org/10.1086/648600 Published online by Cambridge University Press

is the reduced density matrix (here tr_S and tr_B denote, respectively, the partial traces over the degrees of freedom of the open system S and of the environment B. As the dynamics of the total system is Hamiltonian, the total density matrix evolves unitarily, hence the time development of the reduced density matrix may be represented in the form

$$\rho_S(t) = tr_B\{U(t,0)\rho(0)U^\dagger(t,0)\}, \tag{A4}$$

where the initial state of the total system at time $t = 0$ is given by $\rho(0)$ and $U(t,0)$ is the time-evolution operator of the total system over the time interval from $t = 0$ to $t > 0$. The corresponding differential form of the evolution is obtained from a partial trace over the environment of the von Neumann equation,

$$\frac{d}{dt}\rho_S(t) = -itr_B[H(t), \rho(t)]. \tag{A5}$$

Given that the initial state of the total system is of the form $\rho(0) = \rho_S(0) \otimes \rho_B$, the dynamics expressed through (A4) can be viewed as a map of the state space of the reduced system which maps the initial state $\rho_S(0)$ to the state $\rho_S(t)$ at time $t \geq 0$,

$$\rho_S(0) \to \rho_S(t) = V(t)\rho_S(0). \tag{A6}$$

For a fixed t this map is known as a 'dynamical map'. Considered as a function of t it provides a one-parameter family of dynamical maps. If the characteristic time scale over which the reservoir correlations functions decay are much smaller than the characteristic time scale of the system's relaxation, it is justified to neglect memory effects in the reduced system dynamics and to expect a Markovian type behavior, which may be formalized by the semigroup property

$$V(t_1)V(t_2) = V(t_1) + V(t_2), \quad t_1, t_2 \geq 0. \tag{A7}$$

The one-parameter family of the dynamical maps then becomes a quantum dynamical semigroup. Introducing the corresponding generator $\mathcal{L}$ one immediately obtains an equation of motion for the reduced density matrix of the open system S of the form

$$\frac{d}{dt}\rho_S(t) = \mathcal{L}\rho_S(t). \tag{A8}$$

Such an equation is called a Markovian quantum master equation (MME). The most general form of the generator $\mathcal{L}$ is provided by a

https://doi.org/10.1086/648600 Published online by Cambridge University Press

number of theorems, for example, Lindblad's (1976), according to which

$$\mathcal{L}\rho_S = -i[H, \rho_S] + \sum_i \gamma_i \left(A_i \rho_S A_i^\dagger - \frac{1}{2} A_i^\dagger A_i \rho_S - \frac{1}{2} \rho_S A_i^\dagger A_i \right). \quad \text{(A9)}$$

Here H is the generator of the coherent part of the evolution (which need not be identical to the system's Hamiltonian) and the A_i are system's operators with corresponding relaxations times γ_i. In a number of physical situations a quantum master equation whose generator is exactly of the form (A8), known as the Lindblad form, can be derived from the underlying microscopic theory under certain approximations. Two of the most important cases are the singular coupling limit and the weak coupling limit.

The SCL is suitable when one wishes to describe an open quantum system strongly driven by some external macroscopic device with asymptotic states far from equilibrium. To accelerate the decay of the reservoir's correlation one then rescales H_B and H_I. In this limit there is no restriction on the time dependence of H_S, but the price paid is infinite temperature for the bath. The WCL is much more physical and is obtained for all temperatures, but it requires slower system's dynamics. Both derivations must, of course, be consistent with the constraints of thermodynamics, which in this context are formalized as the KMS condition (see, e.g., Kastler 1978), namely, that the reservoir is considered as a heat bath, that is, in thermal equilibrium, so that its entropy and temperature could be defined, and that the total system (system + bath) relax to the equilibrium state imposed by these constraints.

REFERENCES

Aaronson, Scott (2004), "Multilinear Formulas and Skepticism of Quantum Computing", *STOC '04: Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*. New York: Association for Computing Machinery, 118–127.

Aharonov, Dorit (1998), "Quantum Computation", in Dietrich Stauffer (ed.), *Annual Reviews of Computational Physics*, vol. 6. Singapore: World Scientific, 259–346.

Aharonov, Dorit, and Michael Ben-Or (1996), "Fault Tolerant Quantum Computation with Constant Error", http://arxiv.org/abs/quant-ph/9611025.

Aharonov, Dorit, Michael Ben-Or, Russell Impagliazzo, and Noam Nisan (1996), "Limitations of Noisy Reversible Computation", http://arxiv.org/abs/quant-ph/9611028v1.

Aharonov, Dorit, Alexei Kitaev, and John Preskill (2006), "Fault-Tolerant Quantum Computation with Long-Range Correlated Noise", *Physical Review Letters* 96: 050504.

Alicki, Robert (2004), "Comments on 'Fault-Tolerant Quantum Computation for Local Non-Markovian Noise'", http://arxiv.org/abs/quant-ph/0402139v1.

——— (2007) "Comment on 'Resilient Quantum Computation in Correlated Environments: A Quantum Phase Transition Perspective' and 'Fault-Tolerant Quantum Computation with Long-Range Correlated Noise'", http://arxiv.org/abs/quant-ph/0702050v1.

Alicki, Robert, and Mark Fannes (2001), *Quantum Dynamical Systems*. Oxford: Oxford University Press.

Alicki, Robert, Michal Horodecky, Pawel Horodecky, and Ryszard Horodecky (2002), "Dy-

https://doi.org/10.1086/648600 Published online by Cambridge University Press

namical Description of Quantum Computing: Generic Nonlocality of Quantum Noise", *Physical Review* A 65: 062101.
Alicki, Robert, Daniel Lidar, and Paulo Zanardi (2006), "Internal Consistency of Fault-Tolerant Quantum Error Correction in Light of Rigorous Derivations of the Quantum Markovian Limit", *Physical Review* A 73: 052311.
Barenco, Adriano, Andre Berthiaume, David Deutsch, Artur Ekert, Richard Jozsa, and Chiara Macchiavello (1996), "Stabilisation of Quantum Computations by Symmetrisation", http://arxiv.org/abs/quant-ph/9604028.
Bergman, Peter, and Joel Lebowitz (1955), "New Approach to Non-equilibrium Processes", *Physical Review* 99 (2): 578–587.
Berry, Michael (1978), "Regular and Irregular Motion", in S. Jorna (ed.), *Topics in Nonlinear Dynamics*. New York: American Institute of Physics, 16–120.
Blatt, J. M. (1959), "An Alternative Approach to the Ergodic Problem", *Progress in Theoretical Physics* 22 (6): 745–756.
Boltzmann, Ludwig ([1895] 1964), *Lectures on Gas Theory*. Reprint. Translated by Steven G. Brush. Berkeley: University of California Press.
Bonetto, Federico, Joel Lebowitz, and L. Rey-Bellet (2000), "Fourier's Law: A Challenge to Theorists", in A. Fokas, A. Grigoryan, T. Kibble, and B. Zergarlinski (eds.), *Mathematical Physics 2000*. London: Imperial College Press, 128–150.
Borel, Emile (1914), *Le Hasard*. Paris: Alcan.
Breuer, Heinz-Peter, and Francesco Petruccione (2003), "Concepts and Methods in the Theory of Open Quantum Systems", *Lecture Notes in Physics* 622: 65–79.
Brown, Harvey, and Jos Uffink (2001), "The Origins of Time-Asymmetry in Thermodynamics: The Minus First Law", *Studies in History and Philosophy of Modern Physics* 32 (4): 525–538.
Brush, Steven G. (1976), *The Kind of Motion We Call Heat*. Vols. 1–2. Amsterdam: North-Holland.
Bub, Jeff (1997), *Interpreting the Quantum World*. Cambridge: Cambridge University Press.
Burbury, Samuel H (1895), "Boltzmann's Minimum Function", *Nature* 51: 320.
Davies, Brian E. (1974), "Markovian Master Equations", *Communication in Mathematical Physics* 39: 91–110.
——— (1978), "A Model of Heat Conduction", *Journal of Statistical Physics* 18: 161–170.
Dieks, Dennis (1982), "Communication by Electron-Paramagnetic-Resonance Devices", *Physics Letters* 92A: 271.
Dyakonov, Michael (2006), "Is Fault-Tolerant Quantum Computation Really Possible?", http://arxiv.org/abs/quant-ph/0610117.
Dyson, Freeman (2009), "Leaping into the Grand Unknown", *New York Times Review of Books* 56 (6).
Earman, John (1986), "The Problem of Irreversibility", in Arthur Fine and Peter Machamer (eds.), *PSA 1986: Proceedings of the 1986 Biennial Meeting of the Philosophy of Science Association*, vol. 2. East Lansing, MI: Philosophy of Science Association, 226–233.
Earman, John, and Miklos Redei (1996), "Why Ergodic Theory Does Not Explain the Success of Equilibrium Statistical Mechanics", *British Journal of Philosophy of Science* 47: 63–78.
Emerson, Josef, Marcus Silva, Osama Moussa, Colm Ryan, Martin Laforest, Jonathan Baugh, David G. Cory, and Raymond Laflamme (2007), "Symmetrized Characterization of Noisy Quantum Processes", *Science* 317: 1893.
Gaitan, Frank (2008), *Quantum Error Correction and Fault Tolerant Quantum Computing*. Boca Raton, FL: CRC Press.
Gallavotti, Giovanni, and E. G. D Cohen (1996), "Dynamical Ensembles and Stationary States", *Journal of Statistical Physics* 80: 931–970.
Gea-Banacloche, Jean (2002), "Minimum Energy Requirements for Quantum Computation", *Physical Review Letters* 89: 217901.
Gemmer, J., M. Michel, and G. Mahler (2004), *Quantum Thermodynamics: Lecture Notes in Physics*. Vol. 657. Berlin: Springer.
Goldreich, Oded (2005), "On Quantum Computing", http://www.wisdom.weizmann.ac.il/~oded/on-qc.html.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

Goldstein, Shelly, Joel Lebowitz, Roderich Tumulka, and Nino Zanghi (2006), "Canonical Typicality", *Physical Review Letters* 96: 050403.
Gottesman, Danny (1997), "Theory of Fault-Tolerant Quantum Computation", *Physical Review* A 57: 127.
Hagar, Amit (2005), "The Foundations of Statistical Mechanics: Questions and Answers", *Philosophy of Science* 72: 468–478.
Hahn, Erwin (1950), "Spin Echoes", *Physical Review* 80: 580–594.
Haroche, Serge, and Jean-Michel Raimond (1996), "Quantum Computing: Dream or Nightmare?", *Physics Today* 8: 51–52.
Hines, Andrew P., and Philip Stamp (2007), "Decoherence in Quantum Walks and Quantum Computers", http://arxiv.org/abs/0711.1555.
Joos, Erich, H. Dieter Zeh, Claus Kiefer, Domenico Giulini, Joachim Kupsch, and Ion-Olimpiu Stamatescu (1996), *Decoherence and the Appearance of a Classical World in Quantum Theory*. Heidelberg: Springer-Verlag.
Kalai, Gil (2005), "Thoughts on Noise and Quantum Computation", http://arxiv.org/abs/quant-ph/0508095.
——— (2006), "How Quantum Computers Can Fail", http://arxiv.org/abs/quant-ph/0607021v3.
——— (2008), "Detrimental Decoherence", http://arxiv.org/abs/0806.2443.
——— (2009), "Quantum Computers: Noise Propagation and Adversarial Noise Models", http://arxiv.org/abs/0904.3265.
Kastler, Daniel (1978), "Foundations of Equilibrium Quantum Statistical Mechanics", in G. Dell'Antonio, S. Doplicher, and G. Jona-Lasinio (eds.), *Mathematical Problems in Theoretical Physics*. Berlin: Springer, 106–123.
Kaye, Phillip, Ray Laflamme, and Mike Mosca (2007), *An Introduction to Quantum Computing*. Oxford: Oxford University Press.
Kempe, Julia (2005), "Approaches to Quantum Error Correction", *Séminaire Poincaré* 2: 1–29.
Khoon Ng, Hui, and John Preskill (2008), "Fault-Tolerant Quantum Computation versus Gaussian Noise", http://arxiv.org/abs/0810.4953.
Knill, Emanuel (2005), "Quantum Computing with Realistically Noisy Devices", *Nature* 434: 39–44.
——— (2006), "On Protected Realizations of Quantum Information", http://arxiv.org/abs/quant-ph/0603252.
Knill, Emanuel, and Ray Laflamme (1996), "Concatenated Quantum Codes", http://arxiv.org/abs/quant-ph/9608012.
Knill, Emanuel, Ray Laflamme, and Wojciech Zurek (1996), "Accuracy Threshold for Quantum Computation", http://arxiv.org/abs/quant-ph/9610011.
——— (1998), "Resilient Quantum Computation: Error-Models and Thresholds", *Science* 279: 342–345.
Kolmogorov, Andrey (1965), "Three Approaches to the Quantitative Definition of Information", *Problems of Information Transmission* 1 (1): 1–7.
Laflamme, Ray, Cesar Miquel, Juan Pablo Paz, and Wojciech Zurek (1996), "Perfect Quantum Error Correction Code", *Physical Review Letters* 77: 198.
Landauer, Rolf (1995), "Is Quantum Mechanics Useful?", *Philosophical Transactions of the Royal Society: Physical and Engineering Sciences* 1703 (353): 367–376.
Lanford, O. (1981), "The Hard Sphere Gas in the Boltzmann-Grad Limit", *Physica* A 106: 70–76.
Lebowitz, Joel (1978), "Exact Results in Non-equilibrium Statistical Mechanics: Where Do We Stand?", *Progress in Theoretical Physics* (Suppl.) 64: 35–49.
Lenstra, A. K., and H. W. Lenstra Jr. (1990), "Algorithms in Number Theory", in J. van Leeuwen (ed.), *Handbook of Theoretical Computer Science*, vol. A, *Algorithms and Complexity*. Amsterdam: Elsevier, 673–715.
Levin, Leonid (2003), "The Tale of One-Way Functions", *Problems in Information Transmission* 39: 92–103.
Lindblad, Goran (1976), "On the Generators of Quantum Dynamical Semigroups", *Communications in Mathematical Physics* 48: 119–130.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

MacWilliams, F. J., and N. J. A. Sloane (1977), *The Theory of Error-Correcting Codes*. New York: North-Holland.

Margolus, Norman, and Lev Levitin (1998), "The Maximum Speed of Dynamical Evolution", *Physica* D 120: 188–195.

Mermin, David (2007), *Quantum Computer Science: An Introduction*. Cambridge: Cambridge University Press.

Nielsen, Michael, and Isaac Chuang (2000), *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.

Popescu, Sandu, Anthony J. Short, and Andreas Winter (2006), "Entanglement and the Foundations of Statistical Mechanics", *Nature Physics* 2: 754–758.

Preskill, John (1998), "Reliable Quantum Computers", *Proceedings of the Royal Society of London* A 454: 385–410.

——— (2007), "Fault-Tolerant Quantum Computation against Realistic Noise". Talk given at the First International Conference on Quantum Error Correction, University of Southern California, December, http://qserver.usc.edu/qec07/QEC07/JohnPreskill.pdf.

Price, Huw (2003), "Burbury's Last Case: The Mystery of the Entropic Arrow", in Craig Callender (ed.) *Time, Reality and Experience*. Cambridge: Cambridge University Press, 19–56.

Ridderboss, Katinka, and Michael Redhead (1998), "The Spin Echo Experiments and the Second Law of Thermodynamics", *Foundations of Physics* 28 (8): 1237–1270.

Rivest, Ronald, Adi Shamir, and Leonid Adleman (1978), "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems", *Communications of the ACM* 21 (2): 120–126.

Rovelli, Carlo (1994), "Comment on Protective Measurement", *Physical Review* A 50: 2788–2793.

Santayana, George (1905), *The Life of Reason; or, The Phases of Human Progress*, vol. 1, *Reason in Common Sense*. New York: Scribner's.

Shenker, Orly (2000), "Interventionism in Statistical Mechanics—Some Philosophical Remarks", University of Pittsburgh PhilSci Archive, http://philsci-archive.pitt.edu/archive/00000151.

Shor, Peter (1994), "Algorithms for Quantum Computation: Discrete Log and Factoring", in S. Goldwasser (ed.), *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science*. Los Alamitos, CA: IEEE Computer Society, 124–134.

——— (1995), "Scheme for Reducing Decoherence in Quantum Computer Memory", *Physical Review* A 52: 2493–2496.

——— (1996), "Fault-Tolerant Quantum Computation", http://arxiv.org/abs/quantph/9605011.

Simon, Daniel R. (1994), "On the Power of Quantum Computation", in S. Goldwasser (ed.), *Proceedings of the 35th Symposium on Foundations of Computer Science*. Los Alamitos, CA: IEEE Computer Society, 116–123.

Steane, Andrew (1996), "Error Correcting Codes in Quantum Theory", *Physical Review Letters* 77: 793.

Tasaki, Hal (1998), "From Quantum Dynamics to Canonical Distribution: General Picture and Rigorous Example", *Physical Review Letters* 80: 1373–1376.

Terhal, Barbara, and Guido Burkard (2005), "Fault-Tolerant Quantum Computation for Local Non-Markovian Noise", *Physical Review* A 71: 012336.

Uffink, Jos (2001), "Bluff Your Way in the Second Law of Thermodynamics", *Studies In History and Philosophy of Modern Physics* 32: 305–394.

Unruh, William (1995), "Maintaining Coherence in Quantum Computers", *Physical Review* A 51: 992–997.

von Neumann, John (1956), "Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components", in W. Ross Ashby, J. McCarthy and C E. Shannon (eds.) *Automata Studies*. Princeton, NJ: Princeton University Press, 43–98.

Wolfram, Stephen (2001), *A New Kind of Science*. Champaign, IL: Wolfram Media.

https://doi.org/10.1086/648600 Published online by Cambridge University Press

Wootters, William, and Wojciech Zurek (1982), "Single Quantum Cannot Be Cloned", *Nature* 299: 802.
Zanardi, Paulo, and Mario Rasetti (1997), "Noiseless Quantum Codes", *Physical Review Letters* 79: 3306.

https://doi.org/10.1086/648600 Published online by Cambridge University Press