Three Proofs of a Constructive Commuting Quantum Lovász Local Lemma

Toby S. Cubitt,[1] Martin Schwarz,[2] Frank Verstraete,[2] Or Sattath,[3] and Itai Arad[3]

[1] *Departamento de Análisis Matemático, Universidad Complutense de Madrid, Spain*
[2] *Faculty of Physics, University of Vienna, Boltzmanngasse 7, A-1090 Vienna, Austria*
[3] *Department of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel*

(Dated: August 30, 2011)

The recently proven Quantum Lovász Local Lemma generalises the well-known Lovász Local Lemma. It states that, if a collection of subspace constraints are "weakly dependent", there necessarily exists a state satisfying all constraints. It implies e.g. that certain instances of the quantum k–QSAT satisfiability problem are necessarily satisfiable, or that many-body systems with "not too many" interactions are never frustrated. However, the QLLL only asserts existence; it says nothing about how to *find* the state. Inspired by Moser's breakthrough classical results, we present a constructive version of the QLLL in the setting of commuting constraints, proving that a simple quantum algorithm converges efficiently to the required state. In fact, we provide three different proofs, all of which are independent of the original QLLL proof. So these results also provide independent, constructive proofs of the commuting QLLL itself, but strengthen it significantly by giving an efficient algorithm for finding the state whose existence is asserted by the QLLL.

The Lovász Local Lemma (LLL), proven by Erdös and Lovász [1], is a well-known and widely-used result in probability theory. It states that if individual events are "not too" dependent on each other and occur with "not too high" a probability, then there is a positive probability that none of them occur. (This is a non-trivial extension of the trivial fact that, if the individual events were completely independent, and if none of them occurred with certainty, then there would be some probability that none of them occur.) Applications of the LLL abound. It is frequently invoked in order to prove existence of some mathematical object via the probabilistic method. For example, it can be used to prove existence of solutions to instances of boolean satisfiability problems such as k–SAT.

Recently, Ambainis, Kempe and Sattath [2] succeeded in generalising the LLL to the quantum setting. Viewed from one perspective, the Quantum Lovász Local Lemma (QLLL) has little to do with quantum physics; rather, it is a significant mathematical generalisation of the standard LLL to a geometrical setting. However, viewed from another perspective, the quantum version is closely related to current topics of physics research. For example, just as the LLL can be applied to k–SAT problems, the QLLL can be applied to its quantum generalisation, k–QSAT [3]. One way of expressing the k–QSAT problem is: does the Hamiltonian $H = \sum_i \Pi_i$ have a zero-energy ground state, where Π_i are positive-semidefinite local interaction terms (which can without loss of generality be taken to be projectors). This is equivalent to asking whether the ground state of a local Hamiltonian is *frustrated*, an important topic in many-body physics. The QLLL implies that many-body systems in which each particle interacts with "not too many" others are *never* frustrated.

The LLL and its quantum counterpart assert existence, e.g. of a satisfying assignment to a k–SAT instance. But they give no indication as to how to *find* this assignment. In a breakthrough result, Moser [4, 5] recently gave a beautiful proof of a constructive version of the classical LLL. This not only gives an independent proof of the LLL, but does so by giving an efficient algorithm for *finding* e.g. the satisfying assignment to the k–SAT instance. Furthermore, the algorithm is the simplest imaginable: in the k–SAT context, it involves repeatedly picking a clause at random, checking if it's satisfied by the current assignment, and if not, resampling the variables involved in that clause at random. Though there appears on the face of it to be no reason why this algorithm should terminate, let alone find the right state, Moser proved that the expected time until this simple procedure finds the desired state in fact scales linearly!

The conjunction of these two results, the quantum generalisation to the QLLL and Moser's constructive classical LLL, poses a natural question: is there a constructive version of the quantum Lovász Local Lemma? A constructive version of the QLLL is particularly interesting, as it would provide an efficient algorithm for preparing ground states of a large class of many-body quantum Hamiltonians on a quantum computer. Such a procedure has obvious applications in quantum simulation of condensed-matter systems on the physics side, and to quantum satisfiability problems on the computer science side. Moreover, just as the quantum Fourier transform can be viewed as a subroutine for preparing interesting entangled states, a constructive QLLL would provide an efficient subroutine for preparing an entirely new class of entangled states. One might optimistically hope that a constructive QLLL could point the way towards new quantum algorithms.

Our main result is to prove a constructive version of the *commuting* quantum Lovász Local Lemma:

Theorem 1 (Constructive Commuting QLLL)**.** *Let $\Pi_1, \Pi_2, \ldots, \Pi_m$ be commuting projectors that act on arbitrary subsets of n qudits, let $\Gamma(\Pi_i)$ denote the set of projectors which act on a qudit in common with Π_i, and define the relative dimension $R(\Pi_i) = \operatorname{rank}(\Pi_i)/\dim(\mathcal{H})$ (where $\mathcal{H}$ is the Hilbert space on which Π_i acts). If there* *exist values $0 \le x_1, x_2, \ldots, x_m \le 1$ such that*

$$R(\Pi_i) \le x_i \cdot \prod_{\Pi_j \in \Gamma(\Pi_i)} (1 - x_j), \tag{1}$$

then there exists a joint state ρ of the qudits such that $\forall i : \mathrm{Tr}[\Pi_i \rho] = 0$. Moreover, there is a quantum algorithm that converges to a state ρ such that $\mathrm{Tr}[\Pi_i \rho] \le \varepsilon$ in time

$$O\biggl(n + \frac{m}{\varepsilon} \sum_{i=1}^{m} \frac{x_i}{1 - x_i} \cdot |\Pi_i|\biggr), \tag{2}$$

where $|\Pi_i|$ is the number of qudits on which Π_i acts.

There are two significant challenges in generalising Moser's constructive LLL to the quantum setting. First, the state we are trying to construct may be highly entangled, whereas the algorithm only has access to measurements of the local projectors Π_i defining the problem. Another way of expressing this in terms of the Hamiltonian H defining a k–QSAT instance is that, in the classical setting, we know in advance in which basis the overall Hamiltonian H is diagonal—the computational basis—and the local projectors Π_i are local in this same basis. In the quantum setting, the basis which diagonalises the overall Hamiltonian is only defined globally and can be highly entangled, and the projectors Π_i will not in general be local in this basis. Note that this challenge remains just as problematic even if the projectors Π_i commute. In the commuting case of the QLLL, we may know a priori that there exists a basis which diagonalizes all projectors simultaneously, but this basis is still only defined globally and the ground state can still be highly entangled. (Stabilizer states [6] give a simple example of commuting Hamiltonians with highly entangled ground states.) This and related questions in the commutative setting have recently gained attention [7–9] in the context of Hamiltonian complexity.

The second challenge comes from non-commutativity: quantum states are disturbed by measurement. The classical algorithm is free to check which k–SAT clauses are currently satisfied, without any impact on the current assignment. But quantum mechanically, even if we measure a k–QSAT projector Π_i to be "satisfied" (i.e. we obtain the desired $\mathbb{1} - \Pi_i$ outcome upon performing the $\{\Pi_i, \mathbb{1} - \Pi_i\}$ measurement), the measurement can disturb the state such that another Π_j that was previously satisfied no longer is.

Here, we address and give a complete solution to the first of these two challenges: we prove a constructive version of the commuting QLLL (i.e. the case in which all the Π_i commute). A priori, it is not at all clear that Moser's proof extends even to the commuting quantum case, for the reasons discussed above: the basis that diagonalises a set of commuting local projectors can be highly entangled. Furthermore, the local projectors Π_i are in general no longer local in that basis. So even if we knew the right basis, running Moser's algorithm in the diagonal basis would result in throwing away the entire state and starting afresh every time we obtain the wrong outcome for any measurement, which cannot possibly give an efficient algorithm.

Nonetheless, by extending the proof techniques of Moser and Tardos in a more subtle way, we *are* able to prove a constructive version of the commuting QLLL. Moreover, the quantum algorithm involved in Theorem 1 is just the natural quantum generalization of [5], and almost the simplest imaginable. It starts with a uniformly random state (i.e. the maximally mixed state), then in each iteration it chooses one of the projectors at random, and measures it. If the projector is violated (i.e. outcome Π_i is obtained when $\{\Pi_i, \mathbb{1} - \Pi_i\}$ is measured), then the state of those qudits on which that projector acts non-trivially is replaced by a fresh uniformly random state. A priori, there appears to be no reason to expect that this apparently naïve algorithm will find the correct state. We prove that it in fact converges efficiently; indeed, the run-time scales only polynomially in the problem parameters and $1/\epsilon$, where ϵ is the desired precision defined in Theorem 1.

Whilst the algorithm is straightforward, its analysis is not. We provide *three* different proofs of the constructive commuting QLLL, using three very different approaches, each of which makes significant technical contributions of its own.

The first proof generalises the probabilistic approach of Moser and Tardos [5], and proves a constructive version of the most general form of the commuting QLLL, as stated in Theorem 1. (As in the classical setting, this includes the k–QSAT, or "symmetric", version as a special case). A key technical contribution of this proof is the replacement of the classical coupling argument used in the [5] proof, by a *quantum coupling argument*, which uses a coupling by entanglement. To our knowledge, this is the first example of a quantum coupling argument, used as a proof technique to establish convergence of a quantum stochastic process, which may be of independent interest. Even though the entanglement is not used directly as a resource by the algorithm, it is the unique properties of entanglement that allow us to prove via the quantum coupling that the algorithm can find the correct global basis even though it only has access to local measurements.

Coupling arguments have proved to be a very powerful technique in probability theory, often providing the simplest or even the only proof of many results [10]. Our quantum generalisation of the coupling method is no exception, providing an elegant and concise proof of Theorem 1. But coupling arguments often look a little like "black magic", and in our case this makes it difficult to gain significant insight or intuition into why the algorithm succeeds. In the second proof, we replace the coupling argument with a combinatorial proof. Whilst (as often happens) this is significantly more involved than the coupling argument, it is much more explicit. It demonstrates how the algorithm can be understood as a quantum stochastic process produced by iterated measurement, and leads to new results on such measurement processes.

Although not presented this way in the published papers, Moser's original constructive LLL became widely known in an information-theoretic version implicit in [4] (and made explicit in his STOC talk and in [11]), which presents the proof as an elegant entropy compression argument. Moser showed that, if we keep track of the history of the algorithm, the entropy of the information required to store the complete history grows slower than the entropy drawn from the random source, used to resample variables. (In particular, every time the algorithm finds a k–SAT clause to be violated, after a constant initial overhead the new history can be stored using slightly fewer bits than the number of new random bits used to resample the variables in the clause. Thus every clause violation allows the stored information to be compressed faster than the rate at which randomness is used.) But by Shannon's noiseless coding theorem, we cannot compress a uniform random source below its entropy, leading to a contradiction unless the algorithm halts successfully with high probability before it can effectively compress below the entropy. Our third proof generalises this entropic argument to the quantum setting. It proves a slightly weaker version of the symmetric case of the commuting QLLL, in which the conditions on the projectors Π_i are slightly strengthened. But it gives interesting information-theoretic insight into how the algorithm works.

The Lovász Local Lemma has found a very wide range of applications in the 35 years since it was first proven. As it is often used to prove existance of a combinatorial object, as part of an application of the probabilistic method, Moser's recent constructive LLL makes all these existence proofs constructive, which allows these combinatorial objects to be studied directly. The newer quantum LLL has already found applications in areas such as Hamiltonian complexity [2]. So it is no surprise that the constructive QLLL has applications to many-body physics, such as providing a new quantum algorithm for cooling to the ground states of certain many-body Hamiltonians.

Mathematically, our constructive QLLL also has fruitful applications e.g. to the study of CP maps, providing a completely new technique for proving convergence rates of quantum Markov processes to their steady states. For example, the natural dissipative state-engineering map

$$\mathcal{E}(\rho) = \frac{1}{m}\sum_i (\mathbb{1}-\Pi_i)\rho(\mathbb{1}-\Pi_i) + \mathrm{Tr}_{/[i]}[\Pi_i\rho]\otimes\frac{\mathbb{1}_{[i]}}{d^k} \quad (3)$$

was introduced by Verstraete, Wolf and Cirac [12], who showed that it eventually converges to the ground state of the Hamiltonian $H=\sum_i h_i$ if this is frustration-free (where h_i are positive-semidefinite local interaction terms with support Π_i). We use our constructive commuting QLLL to show that, if the local terms h_i commute and do not act on "too many" particles in common (i.e. Π_i satisfy the QLLL conditions of Theorem 1), then the Hamiltonian *is* frustration-free, and moreover the map $\mathcal{E}$ gives fast convergence to the ground state, in time polynomial in the number of local terms and the desired precision (which in this case is the energy of the resulting state), *independent* of the number of particles or their local dimension.

We have resolved one of the two challenges in generalising Moser's results to the quantum setting, allowing us to prove a constructive version of the commuting QLLL. To prove a constructive version of the general, non-commutative QLLL of [2] would require overcoming the second challenge: coping with non-commutativity and the consequences of the measurement-disturbance issue in quantum mechanics. By providing three different proofs of our result, which contribute three quite different sets of new tools and techniques, we have opened up a number of approaches to generalising our result and extending it to a constructive version of the full, non-commutative QLLL. (Indeed, modulo a technical conjecture that is supported by numerical evidence, but appears difficult to prove, our result does extend.) The non-commutative case remains an interesting and challenging open problem.

[1] P. Erdös and L. Lovász, Infinite and finite sets **2**, 609 (1975).

[2] A. Ambainis, J. Kempe, and O. Sattath, "A quantum Lovász Local Lemma," arXiv:0911.1696[quant-ph] (2009).

[3] S. Bravyi, "Efficient algorithm for a quantum analogue of 2-SAT," arXiv:quant-ph/0602108 (2006).

[4] R. A. Moser, in *Proceedings of the 41st annual ACM symposium on Theory of computing (STOC)* arXiv:0810.4812[cs.DS] (2009) arXiv:0810.4812[cs.DS].

[5] R. Moser and G. Tardos, Journal of the ACM (JACM) **57**, 1 (2010), arXiv:0903.0544[cs.DS].

[6] M. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

[7] D. Aharonov and L. Eldar, "On the commuting local hamiltonian problem, and tight conditions on topological order," arXiv:1102.0770 (2011).

[8] N. Schuch, "Complexity of commuting hamiltonians on a square lattice of qubits," arXiv:1105.2843[quant-ph] (2011).

[9] S. Bravyi and M. Vyalyi, "Commutative version of the k-local hamiltonian problem and common eigenspace problem," arXiv:quant-ph/0308021 (2003).

[10] T. Lindvall, *Lectures on the Coupling Method* (Dover Publications, 2002).

[11] T. Tao, "Moser's entropy compression argument," `http://terrytao.wordpress.com/2009/08/05/mosers-entropy-compression-argument/` (2009).

[12] F. Verstraete, M. M. Wolf, and J. I. Cirac, Nature Physics **5**, 633 (2009).