Applications of topology to Weyl semimetals and quantum computing

O'Brien, T.E.

Summary

This thesis is split between different topics over two fields. The first two chapters are concerned with studying and predicting electronic properties of an exotic class of new materials — Weyl semimetals. The remaining five chapters study some of the required parts for the control and operation of a future quantum computer — the architecture and software, and the layer of error correction needed in between. A theme that connects these parts is topology; the idea that one may find behaviour in a system that depends on macroscopic rather than microscopic features. In quantum information this leads to the rather counterintuitive prediction that these macroscopic degrees of freedom may be designed to be resilient against noise. This is counterintuitive as quantum phenomena traditionally decohere faster in larger systems, not smaller, allowing them to remain mostly unnoticed until the 20th century. In condensed matter, topology allows for the existence of systems with exotic effective particles: electronic excitations in a Weyl semimetal with momentum near specific Weyl points obey the Weyl equation, and low-energy electronic excitations in superconducting nanowires obey the Majorana equation. Engineering such systems without defect or disorder is a very difficult task, so we cannot expect to realise such quasiparticles via fine-tuning alone. Luckily, topological protection allows us to circumvent these concerns, and even construct systems (such as effective single Weyl cones) that are fundamentally prohibited in nature.

In part one of this thesis, we investigate novel features of Weyl semimetals. A key feature arising from the electronic structure of a metal (or any conductor) is how its properties vary under the application of a magnetic field. This has applications in the electronics and computing industry, but also gives a means to study the electronic structure itself. The celebrated de Haas-van Alphen and Shubnikov-de Haas effects (discovered in Leiden), and their related counterparts, have been used with great success to probe metals for the last fifty years. These effects appear as fluctuations of various material properties (e.g. magnetic susceptibility, specific heat, and electrical resistivity), periodic in one over the strength of an applied magnetic field. In chapter 2, we show how a signature of a broad class of Weyl semimetals (type-II Weyl semimetals) appears in measurements of this type. This signature manifests by electrons tunneling between different

allowed orbitals in a manner forbidden by a semi-classical description, in a momentum-space analogue to the celebrated Klein tunneling effect. Interestingly, we find that this effect has a topological protection separate to that of a Weyl semimetal, that emerges from the inability to prevent two lines running around a Mobius band from crossing at at least one point.

One of the most stunning predictions in Weyl semimetals is the chiral magnetic effect, whereby electrons are accelerated parallel to an applied magnetic field. This is in contradiction to the standard behaviour of charged particles which, following the Lorentz rule, are accelerated perpendicular to an applied magnetic field. Indeed, one may prove that the gauge symmetry of the electron forbids an equilibrium current parallel to an applied magnetic field. Thus, the chiral magnetic effect is not observable in a regular metal in equilibrium. However, the gauge symmetry of the electron is broken in a superconductor, re-opening the possibility of such an observation, which we demonstrate and study in chapter 3. This manifests as a charge current balanced by a flow of zero-charge particles in the opposite direction to balance the heat current (which must remain net-zero). The flow of these charged particles depends on the chirality of an effective single Weyl cone in the bulk of the system, which may be switched by an applied supercurrent, thus providing a 'chirality switch' for this effect.

In part two of this thesis we switch our focus to quantum error correction, where topology is used to enhance control over quantum degrees of freedom, rather than ensuring their existence. Although it is widely-accepted that fault tolerant quantum computing is eventually possible, accurate numerical predictions of the performance of current efforts to demonstrate scalable full quantum error correction have been lacking, in part due to the complexity of the system. In chapter 4 we rectify this, by performing a full-density matrix simulation of a small quantum error correcting code prototype, known as Surface-17, on superconducting quantum hardware. As part of this, we explicitly detail and develop a scheme for measuring the experimental performance of the Surface-17 chip, and provide and benchmark a decoder to analyze the information from repeated parity checks performed during this experiment. The density-matrix simulation toolbox developed during this work allows us to accurately model experimental performance as noise parameters are varied, which has proved essential in both theory and experimental studies since then.

As mentioned in the previous chapter, quantum error correction requires classical processing of repeated syndrome measurements to diagnose and correct local errors before they accumulate to kill the macroscopically-

spread logical quantum information. Optimizing this decoding requires detailed knowledge of the rates at which these local errors occur in order to determine the most-likely set of underlying events from the error syndrome. In chapter 5 we investigate how one may determine this knowledge exactly within the framework of one of the most popular decoding algorithms, based on the minimum-weight perfect matching problem in graph theory. We test and demonstrate this adaptive decoder both in estimating underlying error rates in a static system, and in keeping up with noise fluctuations on the time-scale of a single quantum error correcting experiment. We find that the decoder may keep up with what would be sub-millisecond fluctuations on a superconducting transmon chip, and demonstrate the trade-off between lagging behind said fluctuations and trying too hard to keep up (resulting in insufficient convergence).

The previous two chapters dealt with error correction for the surface code, which is a popular choice for quantum error correction due to the excellent performance of minimum-weight perfect matching decoders. Other topological quantum error correcting codes, such as the color code, have even lower qubit counts and easier access to logical operations than the surface code, however research in these has been somewhat stymied due to the lack of good decoders. In chapter 6, we extend a class of neural network quantum error correction decoders to work on the color code. Importantly, this requires both processing of traditional parity-check measurements, and additional flag measurements, which are required to diagnose a small class of particularly dangerous 'hook' errors that appear in the color code. We demonstrate numerical evidence that our scheme protects against these hook errors, despite using fewer flag measurements than expected in the rigorous theoretical proofs of their performance. This allows us significant savings in measurement overhead, giving us (to the best of our knowledge) the best-performing color code decoder yet. Furthermore, the presented scheme is usable for any quantum error correcting code, opening up a range of interesting possibilities for code design.

In the final part of this thesis, we look towards future implementations of quantum computers and the algorithms to run thereon. Majorana modes, or Majorana bound states, have emerged as an exciting prospect for quantum computing. Their topological protection raises hopes of lower error rates than current state-of-the-art devices (which is potentially the largest barrier to future useful quantum computers, even with quantum error correction). However, most work studying Majorana modes as a platform for quantum computing has paid little attention to their non-abelian quasiparticle nature, save for the purposes of constructing

a universal quantum gate set. This is of importance when looking to perform future quantum chemistry calculations on a quantum computer: these calculations typically study electronic behaviour, the simulation of which on a bosonic qubit system incurs a significant overhead. In chapter 7, we propose a scheme where one might circumvent this overhead on a Majorana architecture by directly combining pairs of Majoranas to make fermions. We demonstrate how one may decompose traditional quantum operations into the language of fermions (albeit in a non-fault tolerant setting), and show a large advantage over traditional quantum computing implementations in simulating lattice fermionic systems.

In order to perform quantum chemistry calculations, or any calculations, on a quantum computer (with a quantum speedup over classical computers), one must tailor new algorithms to quantum devices. The work-horse behind most (if not all) algorithms with a provable quantum speedup is quantum phase estimation. As suggested in the name, this algorithm provides a means to determine the relative phase accumulated by an otherwise stationary state (also known as an eigenstate) as it evolves following the Schrödinger equation (simulated on a quantum register). Traditionally, this has either required accumulating this phase onto multiple ancillary qubits, or precise preparation of the eigenstate itself, both of which come with significant overhead. In the last chapter of this thesis, we demonstrate how one may separate the phase information from a superposition or mixture of eigenstates by single rounds of the quantum phase estimation protocol with a single ancilla qubit. This work pioneers the use of Prony's method to extract phase information from a quantum computer, which is preferable to more traditional Fourier-type methods as it allows for finer resolution of a small number of phases with a low classical computation cost. We finally study the performance of this method, and that of Bayesian post-processing, in the presence of nearby eigenenergies (which make distinction of a single frequency difficult), and two prototypical noise models (including the accurate density-matrix simulations presented in Chapter 4).