Chapter 1

Introduction

In 1994, Peter Shor [45] made a huge discovery in the field of quantum computing. He discovered that a quantum computer can factor integers efficiently, a problem so strongly believed by researchers to be difficult that all major encryption on the internet is based on the inability to factor integers.

Prior to this there was significant work leading up to Shor's discovery. In 1985 Deutsch [19] defined the first complete model of a quantum computer, but it was not for several years that progress was made in understanding the power of the computational model. Deutsch and Jozsa [20] gave the first nontrivial quantum algorithm. Building upon this Bernstein and Vazirani [8] gave the first example of a problem whose quantum algorithm has a super-polynomial speedup over the best classical algorithm. Their main tool is Fourier sampling, i.e. the process of computing the Fourier transform (of some state), followed by a measurement. Simon [46] then gave a problem and a quantum algorithm that has an exponential speedup over the best classical algorithm. It was from these ideas that Shor

developed his algorithms for factoring and discrete log.

Since Simon's and Shor's discoveries, a framework has been established, through which to understand and extend their results. This framework is built on the Hidden Subgroup Problem (HSP): given a function defined on a finite group, constant and distinct on cosets of an unknown subgroup, find a set of generators for the subgroup. A natural generalization of their algorithms solves the HSP over any abelian group, and this fact has become a folklore theorem over time. Shor's algorithms for factoring and discrete log are somewhat more complicated, and the underlying problem may be abstracted as the variant on the HSP where the group is restricted to be abelian, but with the added complication that the sizes of the cyclic factors of the group are unknown. The solution of this problem follows from a general Fourier sampling theorem we prove in this thesis. A different problem was defined and solved by Kitaev [36], called the Abelian Stabilizer Problem, a problem that has factoring and discrete log as special cases. The stabilizer problem seems very general, but reduces to the HSP (with unknown group size). Kitaev's algorithm provides a completely different point of view of how the HSP algorithm works, but the quantum circuit has been shown to be the same as for the HSP.

Another extremely important question is whether or not a quantum computer can efficiently solve NP-Complete problems, a class of hundreds of problems of great practical importance, and widely believed to be intractable for classical computers. Unfortunately, evidence was provided by Bennett, Bernstein, Brassard, and Vazirani [7] that this class cannot be solved efficiently on a quantum computer, either.

There are at least a few known problems that are probably not NP-Complete,

but appear to be hard classically, and these motivate a search for more efficient quantum algorithms. One example is graph isomorphism, a problem important for practical and theoretical reasons [37]. Another example is the Unique Shortest Lattice Vector Problem, which, like factoring, has cryptography systems [1] based on the fact that no efficient algorithm has been found.

This thesis studies the generalization of Shor's algorithms to general abelian groups, the generalization of Simon's algorithm to nonabelian groups, and solves a new problem whose quantum algorithm differs from the HSP framework.

Implicit in Shor's proof is a robustness of Fourier sampling, for certain restricted initial quantum states, with respect to certain changes in the underlying group. Our first result [31] is a generalization of this: we show that for any quantum state, Fourier sampling is robust under the changes in the underlying group. Underlying our proof is the fact that for a fixed vector, the Fourier transform over $\mathbb{Z}_p$ can be obtained by considering it as a over $\mathbb{Z}$, taking the inverse Fourier transform with respect to the circle group to get a continuous function, and then taking p evenly spaced points. This solves the principal difficulty in applying the HSP algorithm to solving actual problems where the underlying group is unknown. The robustness theorem may be thought of as providing a compiler where if a problem can be reduced to the HSP where the group is known, then the algorithm will work even if it is not known.

We then combine these ideas with those of [36], [14] and [35] and construct a new Fourier transform algorithm for any abelian group [32]. This algorithm is faster than the previous one [36] and is extremely simple. Due to the simplicity we are able to improve

our Fourier sampling theorem to work on any abelian group, and we get an especially clean explanation of the cyclic case.

By contrast, the nonabelian HSP is wide open. It includes as a special case the longstanding open question of graph isomorphism, where the group is the symmetric group, S_n. It is natural to carry over the abelian HSP algorithm to the nonabelian case. The algorithm must efficiently compute the Fourier transform, but in the nonabelian case whether this can be done or not must be solved on a group by group basis. For groups we are interested in, such as the symmetric group, it is known [4] how to efficiently compute the quantum Fourier transform. We are mainly interested in how much information can be obtained by Fourier sampling. Ettinger and Høyer [24] showed that the HSP over the dihedral group has polynomial query complexity, but has exponential post-processing time. In this thesis [33] we provide the first classification of how the algorithm for the abelian case of the Hidden Subgroup Problem generalizes for an arbitrary finite nonabelian group. We show that in the case that the hidden subgroup is normal, this algorithm succeeds in reconstructing the subgroup in polynomial time. On the other hand, we give evidence that this algorithm is inadequate to even distinguish a trivial subgroup from an involution in the case of the symmetric group. Independently, Grigni, Schulman, and Vazirani [28] also showed that the algorithm will not work with the reduction, and also showed that an even more general version of the algorithm will not work.

Finally, we give a new quantum algorithm [49] that does not appear to fit into the framework of the HSP. It solves the Shifted Legendre Symbol Problem and its variants. There is some evidence that this is an intractable problem classically, and a closely related

problem has been proposed as a cryptographic primitive. In addition to the fact that the base algorithm is new, there are two other interesting issues that come up. In one variant of the problem a periodic function on a ring is given and the period must be found first. This does not appear to fit into the framework of previous period finding algorithms (namely cyclic subgroups in the HSP problem), and uses our Fourier sampling theorem. The other interesting issue is in the variant defined over finite fields, where the Fourier transform we use must be chosen carefully so that it behaves with respect to both addition and multiplication in the field.

There are other important areas of quantum computation which we do not discuss here. They include quantum error correction, quantum lower bounds, and quantum communication and cryptography. For more on these areas see Nielson and Chuang [41] or Preskill [43].

Chapter 2 provides a reference to notation of objects used throughout the thesis. Chapter 3 defines the model of a quantum computer. Chapter 4 defines the notion of Fourier sampling, and gives the HSP algorithm. In Chapter 5 we show how to relate Fourier transforms over different groups, resulting in a Fourier sampling theorem and a new Fourier transform algorithm. In Chapter 6 we provide the new results on the nonabelian hidden subgroup problem. Chapter 7 solves the Shifted Legendre Symbol Problem.