Quantum search with advice

Ashley Montanaro

Department of Computer Science, University of Bristol, Bristol, UK

Abstract. We consider the problem of search of an unstructured list for a marked element x, when one is given advice as to where x might be located, in the form of a probability distribution. The goal is to minimise the expected number of queries to the list made to find x, with respect to this distribution. We present a quantum algorithm which solves this problem using an optimal number of queries, up to a constant factor. For some distributions on the input, such as certain power law distributions, the algorithm can achieve exponential speed-ups over the best possible classical algorithm. We also give an efficient quantum algorithm for a variant of this task where the distribution is not known in advance, but must be queried at an additional cost. The algorithms are based on the use of Grover's quantum search algorithm and amplitude amplification as subroutines.

Grover's algorithm for search of an unstructured list is one of the greatest successes of the nascent field of quantum computation. However, it is rarely necessary to search the type of databases that we encounter in real life in a completely unstructured fashion. Instead, there is often some prior information about the location of the sought ("marked") item x, which can be used to guide the search. We can formalise this intuition by considering a search problem where the searcher is given access to a probability distribution μ, which hints where the marked item is likely to be. This problem can be stated formally as follows.

- **Problem:** Search with Advice
- **Input:** A function $f : \{1, \dots, n\} \rightarrow \{0, 1\}$ that takes the value 1 on precisely one input x, and an "advice" probability distribution $\mu = (p_y)$, $y \in \{1, \dots, n\}$, where p_y is the probability that $f(y) = 1$.
- **Output:** The marked element x.

It is clear that knowledge of μ can enable a classical algorithm to achieve a significant reduction in the average number of queries to f (with respect to μ) required to find the marked element x. This paper is concerned with the development of *quantum* algorithms for the Search with Advice problem which also use μ, and which obtain significant speed-ups over any possible classical algorithm.

In the main model we study in this paper, which we call the *known* model, μ is known completely beforehand, and can be used to help design an algorithm to find the marked element. We also discuss a second model – the *unknown* model – where μ is not known before the algorithm starts. In both cases, we are interested in the *average* number of queries with respect to μ required to find the marked element, rather than the worst-case number of queries. Previous work has shown that, if one considers the worst-case number of queries to the input required to compute any total function, there can only be at most a polynomial separation between quantum and classical computation. Considering the average number of queries required (over the input) allows one to sidestep these results and hope to obtain *exponential* speed-ups.

The main results of this paper are as follows. First, we give a quantum algorithm for SEARCH WITH ADVICE which is optimal up to constant factors. Assuming without loss of generality that the probability distribution $\mu = (p_x)$ is given in non-increasing order, the algorithm uses an expected number of queries to f which is of the order of $\sum_{x=1}^{n} p_x \sqrt{x}$, which should be compared with the optimal classical expected number of queries, $\sum_{x=1}^{n} p_x\, x$. For certain probability distributions, this represents an exponential (or even super-exponential) improvement in the expected number of queries used. The quantum algorithm is based on the use of an exact variant of Grover's algorithm as a subroutine. Known lower bounds on the query complexity of quantum search are extended to show that this algorithm is optimal, up to constant factors, for any probability distribution μ.

This result is applied to the natural class of power law distributions $p_x \propto x^k$, for some constant $k < 0$. We show that for certain values of k, quantum algorithms deliver very significant reductions in the average number of queries used. In particular, when $-2 < k < -3/2$, a super-exponential separation between quantum and classical computation is obtained in the known model ($O(1)$ vs. $\Omega(n^{k+2})$).

The paper closes by briefly discussing results in the "unknown" model, where information about the probability distribution μ must be obtained by queries at unit cost. Here, we have a quantum algorithm that uses amplitude amplification to achieve an expected number of queries of the order of

$$\left(\sum_{x, p_x > 1/n} \sqrt{p_x} \right) + \sqrt{n} \left(\sum_{x, p_x \leq 1/n} p_x \right).$$