

Quantum search of partially ordered sets

Ashley Montanaro*

Department of Computer Science, University of Bristol,
Woodland Road, Bristol, BS8 1UB, UK.

March 2, 2007

Abstract

We investigate the generalisation of quantum search of unstructured and totally ordered sets to search of partially ordered sets (posets). Two models for poset search are considered. In both models, we show that quantum algorithms can achieve at most a quadratic improvement in query complexity over classical algorithms, up to logarithmic factors; we also give quantum algorithms that almost achieve this optimal reduction in complexity. In one model, we give an improved quantum algorithm for searching forest-like posets; in the other, we give an optimal $O(\sqrt{m})$ -query quantum algorithm for searching posets derived from $m \times m$ arrays sorted by rows and columns. This leads to a quantum algorithm that finds the intersection of two sorted lists of n integers in $O(\sqrt{n})$ time, which is optimal.

1 Introduction

Searching for an object in a set of objects that obey some structure is a fundamental task in computer science. The archetypal example of such a task is finding an integer in a sorted list containing n elements; in this case, binary search can find the marked integer in $O(\log n)$ steps. At the other extreme, any (classical) search algorithm requires $\Omega(n)$ steps to search a completely unsorted n -element list. It is thus of interest to find a framework for search problems that encompasses both of these structures, and interpolates between them.

One approach is to consider the task of searching a partially ordered set (*poset*). Recall that a partial order on a set S is a relation \leq such that, for $a, b, c \in S$, $a \leq a$, $(a \leq b) \wedge (b \leq a) \Rightarrow a = b$, and $(a \leq b) \wedge (b \leq c) \Rightarrow a \leq c$. We define the relation $<$ in the obvious way: $(a < b) \Leftrightarrow (a \leq b) \wedge (a \neq b)$. For any two elements a, b , either $a \leq b$, $b \leq a$, or a and b are incomparable, $a \not\leq b$. We say that a set is totally ordered if none of its elements are incomparable, and unstructured if all of its elements are incomparable.

There are two natural ways to model poset search. In the first model (introduced by Linial and Saks [20], and called the *concrete* model here), we consider the partial order on S

*montanar@cs.bris.ac.uk

to represent constraints on the structure of an unknown *totally* ordered set, identified with the integers. That is, each element $s \in S$ stores an integer $x = S[s]$, which is returned by a query to the element s . The constraint following from the partial order on S is that if $s \leq t$ for some $s, t \in S$, then $S[s] \leq S[t]$. The goal is to find the location at which a (known) arbitrary integer a is stored, or to output that a is not stored in S , using the minimum number of queries to elements of S . We will usually assume that the integers stored in S are all distinct.

Alternatively, in the second model (introduced by Ben-Asher, Farchi and Newman [7], and called the *abstract* model here), the goal is to search for an unknown “marked” element $a \in S$, using the minimum number of queries to an oracle, which operates in the following way. On input of an element $x \in S$, the oracle returns one of $\{<, =, \not\leq\}$. The first two possibilities are returned when $a < x$ and $a = x$, respectively, and the third is returned when either $x < a$ or x and a are incomparable.

We sometimes mention an extension of the search problem to a scenario where multiple different answers are permissible. This extension is different for the two models: in the abstract model, we consider there to be multiple marked elements in the set to be searched, with the goal being to output any of these elements. In the concrete model, the analogous scenario is allowing the possibility for the set to store duplicate integers, i.e. allowing there to exist $s, t \neq s$ such that $S[s] = S[t]$.

To sum up, in the concrete model we know what we are searching for, but not where to find it; in the abstract model, we do not know what we are searching for, but we can perform powerful queries that narrow down the search space to find it.

This paper is concerned with quantum search of posets in both of these models, and in particular with minimising the number of queries to the set required to find the desired element. It is well-known that Grover’s algorithm [16] can find the marked element in an unstructured n -element set using $O(\sqrt{n})$ quantum queries, thus gaining a quadratic advantage over classical computation, and that this reduction is optimal. However, no advantage beyond a constant factor may be achieved for quantum search of a totally ordered set [18].

We then have several questions, motivated by these two examples. Can we find interesting quantum algorithms for search of general posets? Could a reduction in queries of more than the quadratic factor given by Grover’s algorithm be achieved by such an algorithm, or even an exponential reduction? And what is the structure (or otherwise) of the posets for which a quantum computer can gain an asymptotic advantage over classical computation?

1.1 New results

Our first result is that, in both the abstract and concrete models, quantum algorithms can achieve no more than a quadratic reduction (up to a logarithmic factor) in the number of oracle queries to find a marked element. The lower bounds in the two models seem to need different proof techniques: the bound in the abstract model follows from a reduction to the oracle identification problem of Ambainis et al [4], whereas we use structural properties of posets to derive the lower bound in the concrete model.

We give general upper bounds that match these lower bounds up to logarithmic factors. In the abstract model, the upper bound follows from an algorithm of Atici and Servedio [6]. In the concrete model, we give a new and almost optimal quantum algorithm that follows from Dilworth’s Theorem [15] on the decomposition of posets into ordered components.

These general results can be summarised as the following theorem.

Theorem 1.1. *Let S be an n -element poset, and let $D(S)$ and $Q_2(S)$ be the number of queries required for an exact classical or bounded-error quantum (respectively) algorithm to find the marked element in S , in either of the two models discussed above. Then*

$$\begin{aligned}
 D(S) &= O(Q_2(S)^2 \log n) \\
 Q_2(S) &= \begin{cases} O(\sqrt{D(S)} \log n \sqrt{\log \log n}) & (\text{abstract model}) \\ O(\sqrt{D(S)} \log n) & (\text{concrete model}) \end{cases}
 \end{aligned}$$

In both models, we give explicit quantum algorithms for searching specific poset structures. In the abstract model, we give a simple (and nearly optimal) algorithm for searching a class of forest-like posets. For an unstructured set, the algorithm reduces to Grover search, whereas for a totally ordered set it reduces to binary search.

In the concrete model, we give an asymptotically optimal quantum algorithm for searching posets that are derived from 2-dimensional arrays of distinct integers sorted by rows and columns. This gives rise to an optimal quantum algorithm for an apparently unrelated problem: finding the intersection of two sorted lists. Given two lists of at most n integers in increasing order, the algorithm can find an element that appears in both lists in $O(\sqrt{n})$ time, improving on a previous algorithm of Buhrman et al [10], which achieved a time complexity of $O(\sqrt{nc}^{\log^* n})$ for some constant c .

1.2 Previous work

Classically, the question of searching partially ordered sets seems to have first been considered by Linial and Saks [20, 21], who characterised the query complexity of searching posets in their concrete model. They showed that this complexity depends solely (up to constant factors) on the number of ideals of the poset, where an ideal of S is a subset $T \subseteq S$ such that $(x \in T) \wedge (y < x) \Rightarrow (y \in T)$. In particular, they give lower and upper bounds on the complexity of searching for a marked element in an array sorted by rows and columns, and the d -dimensional generalisation thereof.

Ben-Asher, Farchi and Newman [7] introduced the abstract model, and gave an algorithm to find the optimal search strategy in this model for a class of tree-like posets. In this model, it is interesting to note that the problem of determining an optimal search strategy for arbitrary posets is NP-hard, whereas the same question restricted to trees is soluble in polynomial time [12]. In fact, Onak and Parys have recently obtained an $O(n^3 \log n)$ -time algorithm for finding this strategy [23], and also point out that this model is similar to a model of search in graphs, where one queries an edge and is returned the closest endpoint of that edge to the marked element. It was already known that near-optimal search strategies for almost all posets can be produced efficiently [12].

In the case of quantum search, tight upper and lower bounds on query complexity are known for search of unstructured sets [16, 9, 25]. An asymptotically tight lower bound is known for search of totally ordered sets [4, 18]. We will also make use of related results by Aaronson and Ambainis on spatial quantum search [1].

2 Preliminaries

2.1 Quantum query algorithms

In this work, the measure used of the complexity of searching a poset S is usually the number of queries to S required to find the marked element, or report that none exists, rather than the time required for the search (see Section 5 for a brief discussion of this point).

We will assume familiarity with quantum computation, and will use the standard model of quantum query complexity. In this model, a t -query quantum algorithm is a sequence of unitary transformations $U_t O_a U_{t-1} O_a \cdots O_a |\psi\rangle$, where we alternate between “expensive” oracle queries that may depend on an unknown entity a , and “free” arbitrary unitary operations that do not, with the aim being to minimise the number of oracle queries. The oracle O_a is usually taken to be a unitary operator that operates on an n -dimensional input register $|x\rangle$ and d -dimensional output register $|y\rangle$, and encodes an arbitrary function $f_a(x) : \mathbb{Z}_n \mapsto \mathbb{Z}_d$ as follows: $O_a|x\rangle|y\rangle = |x\rangle|y + f_a(x)\rangle$, where addition is taken modulo d .

In the abstract model, we require an oracle $f_a(x)$ that returns something from the set $\{<, =, \not\leq\}$, according to whether the unknown marked element $a < x$, $a = x$ or $a \not\leq x$. However, it will be convenient to instead use a Boolean oracle by adding a parameter $z \in \{0, 1\}$ to give an oracle function $f_a(x, z)$, which acts as follows. $f_a(x, 0) = 1$ if $a \leq x$, and 0 otherwise. $f_a(x, 1) = 1$ if $x = a$, and 0 otherwise. It is clear that a query to $f_a(x)$ is sufficient to simulate a query to $f_a(x, z)$, and querying $f_a(x, 0)$ and $f_a(x, 1)$ is sufficient to simulate $f_a(x)$. The query complexity in the two-parameter model may thus only differ by a factor of at most 2 from the one-parameter model. The model can be extended to allowing more than one marked element in an obvious way, by parametrising the oracle with a set of marked elements A ; then $f_A(x, 0) = 1$ if there exists $a \in A$ with $a \leq x$.

The concrete model is more straightforward; here, the oracle depends only on the integers stored in the set S , and an oracle query to an element x simply returns the integer stored at the element x , i.e. $S[x]$. We usually assume that, for all $x \neq y$, $S[x] \neq S[y]$.

$D(S)$ will denote the worst-case exact classical decision tree complexity of searching for a single marked element in the poset S , and $Q_E(S)$ the equivalent quantum query complexity. $Q_2(S)$ is the quantum query complexity where we are allowed to err with probability $\leq 1/3$ (the “2” refers to 2-sided bounded error). Motivated by binary search, our notion of a poset S that allows “efficient” search is one where the marked element can be found using a number of queries that is polylogarithmic in $|S| = n$. All logarithms will be taken to base 2.

We will make frequent use of an exact variant of Grover’s quantum search algorithm [16].

Theorem 2.1. (Exact Grover search [e.g. [9], [22]])

Let S be an unstructured set of n elements containing either one marked element, or no marked elements. Then there exists an exact quantum algorithm which outputs the marked element, or that no such element exists, using $O(\sqrt{n})$ queries to the set.

2.2 Posets

We will use standard terminology relating to posets. A *chain* in a poset S is a subset $T \subseteq S$, all of whose elements are comparable. Conversely, an *antichain* is a subset whose elements are all incomparable. The *height* $h(S)$ and *width* $w(S)$ of a poset S are the size of the largest chain and antichain in S , respectively. A *subset* of a poset S is a subset of the elements in S that preserves the order relations; conversely, an *extension* of S preserves the elements but may add new order relations. A *section* of S is a subset $T \subseteq S$ such that $(x \in T) \wedge (z \in T) \wedge (x < y < z) \Rightarrow y \in T$. A *maximal element* of S is an element $x \in S$ such that for all $y \in S$, $y \not> x$.

A poset can be represented graphically by its *Hasse diagram*. A Hasse diagram for S is an undirected graph G whose vertices are labelled by the elements of S . We say that b covers a if $b > a$ and there does not exist $c \in S$ such that $a < c < b$. For each pair of vertices a, b , if a covers b then the vertex corresponding to a in the Hasse diagram is connected to, and positioned higher than, that corresponding to b . Figure 1 gives the Hasse diagrams of some example posets.




Figure 1: Hasse diagrams of some posets

A poset S is said to be *tree-like* (forest-like) if its Hasse diagram is a tree (forest) rooted at the maximal element(s) of S .

3 The abstract model

In this section, we consider the problem of searching posets in the model studied by Ben-Asher, Farchi and Newman [7], where a query to an element of a poset S returns information about its relationship to the unknown marked element with respect to the partial order on S .

3.1 Overall relationships

In this model, we can immediately relate quantum and classical search using a reduction to the oracle identification problem, which was originally introduced (in the context of quantum computation) by Ambainis et al [4], and related to computational learning theory by Servedio and Gortler [24]. In this problem, we are given as an oracle an unknown m -bit Boolean function f picked from a known set of functions S , and we must identify f with the minimum number of queries to the oracle (Servedio and Gortler refer to this as *exactly learning* f).

Servedio and Gortler have shown [24] that the quantum and classical query complexities of this task are closely related, and both depend on a parameter which we call γ^{S^1} , which is informally defined as the minimum fraction of the functions in S which a classical algorithm can be certain of removing from consideration with a query to f . To be precise, let S' be a subset of S , and let $S'_{a,b}$ be the subset of those functions in S' that take value b on input a . Then

$$\gamma^S = \min_{S' \subseteq S, |S'| \geq 2} \max_{a \in \{0,1\}^m} \min_{b \in \{0,1\}} \frac{|S'_{a,b}|}{|S'|} \quad (1)$$

The main result of [24] may be stated as:

Theorem 3.1. [24]

Let S be a set of Boolean functions on m bits. Then the quantum query complexity Q of exactly learning a function from S , with a bounded probability of error, obeys the following lower bounds.

$$Q = \Omega\left(\frac{1}{\sqrt{\gamma^S}}\right), \quad Q = \Omega\left(\frac{\log |S|}{m}\right) \quad (2)$$

Also, the deterministic classical query complexity C of the same task obeys the following upper bound.

$$C = O\left(\frac{\log |S|}{\gamma^S}\right) \quad (3)$$

Quantum and classical query complexities are thus related by $C = O(mQ^3)$.

The classical algorithm that achieves this query complexity is quite straightforward, simply consisting of querying the unknown function at the input that, given an adversarial response, reduces the size of the set of remaining possible functions by the largest possible amount.

We now make a connection between the poset search problem and oracle identification. Given a poset, the oracle associated with each possible marked element a is a two-parameter Boolean function $f_a(x, z)$. Distinguishing between these functions is exactly equivalent to finding the hidden a . Thus, in order to find the marked element in an n -element poset, we need to distinguish n Boolean functions on $\lceil \log n + 1 \rceil$ bits. Theorem 3.1 immediately gives the following result.

Theorem 3.2. *Let S be an n -element poset. Then $D(S) = O(\log n Q_2(S)^2)$.*

¹This is Servedio and Gortler's $\hat{\gamma}^C$.

A quadratic reduction in queries is thus the best that can be obtained using a quantum algorithm, up to a logarithmic factor. We now turn to upper bounds on quantum query complexity. There is a straightforward general upper bound of $O(\sqrt{n})$ oracle queries for any poset. This can be seen by noting that, if the oracle $f_a(x, z)$ is queried only with $z = 1$, the problem reduces to unstructured search, so Grover’s algorithm [16] can be used.

Less trivially, Atici and Servedio [6] have given a quantum algorithm for exact learning that can be seen as an analogue of the classical algorithm mentioned in Theorem 3.1. This algorithm immediately applies to poset search, and moreover is efficient (runs in time polynomial in n).

Theorem 3.3. [6] *Let S be an n -element poset. Then*

$$Q_2(S) = O\left(\frac{\log n \log \log n}{\sqrt{\gamma^S}}\right) \quad (4)$$

This upper bound can actually be improved to $Q_2(S) = O\left(\log n \sqrt{\log \log n} / \sqrt{\gamma^S}\right)$. The reason is that the $O(\log \log n)$ factor in Atici and Servedio’s algorithm comes from performing $O(\log \log n)$ rounds of classical probability amplification, which can be replaced by the use of a quantum algorithm of Buhrman et al [11] that performs efficient amplitude amplification to small error probabilities.

In summary, it can be seen that the quantum and classical query complexities of this search problem are completely determined (up to logarithmic factors) by this parameter γ^S . However, it is unclear whether the extension to searching for multiple marked elements has a similar reduction to the oracle identification problem, and whether a suitable adaptation of Atici and Servedio’s algorithm can be applied in this case.

Finally, note that one might consider a more powerful variant of search in this model, where the oracle $f_a(x)$ is extended to return $>$ if the marked element $a > x$ (so the four possible results are “ $<$ ”, “ $=$ ”, “ $>$ ” and “incomparable”). The reduction to the oracle identification problem clearly still holds for this variant, so the results in this section go through unchanged.

3.2 Search in forest-like posets

We say a poset is forest-like if every element in the poset is covered by at most one other element (an example of such a poset is shown in Figure 1). Classically, forest-like posets have proven to be easier to analyse; indeed, algorithms exist [7, 23] for computing the optimal classical decision tree to search these posets in polynomial time, whereas the same problem is NP-hard for general posets [12]. In this section, we present an exact quantum algorithm for searching a forest-like poset S using $O\left(\log n / \sqrt{\gamma^S}\right)$ queries, improving on the previously mentioned bounded-error $O\left(\log n \sqrt{\log \log n} / \sqrt{\gamma^S}\right)$ -query algorithm [6]. Our algorithm improves on Atici and Servedio’s in other ways too: firstly, it reduces to an asymptotically optimal algorithm in the case of search of unstructured and totally ordered sets; secondly,

it can easily be extended to searching for multiple marked elements, with a small penalty in query complexity.

We first consider the case of a single marked element. The principles behind the algorithm that we will describe are very similar to those underlying Atici and Servedio’s. Throughout the algorithm, a subset of possible places that the marked element could be is maintained. We will show that one use of Grover’s algorithm over a set G of size at most $1/\gamma^S$ can be used to reduce the size of this subset by at least half, so $\log n$ repetitions suffice to find the marked element. Crucially, for forest-like posets where there is a single marked element, this use of Grover’s algorithm can be made exact (Theorem 2.1), thus avoiding the need for some number of repetitions to achieve a suitable reduction in the error probability.

The algorithm is explicitly stated as Algorithm 1 below. It uses a function `centralElement` which requires some explanation. Define the weight $wt(v)$ of an element $v \in S$ as $wt(v) = |\{x : (x \in S) \wedge (x \leq v)\}|$. Then `centralElement(S)` returns the element $v \in S$ such that $wt(v)$ is maximised, given that $wt(v) \leq \lceil |S|/2 \rceil$. Such an element will clearly always exist. `siblings(x)` returns the set of elements of S that are covered by the single element that covers x .

Algorithm 1 Search algorithm for forest-like posets

Input: Forest-like poset S containing n elements

Output: Marked element, or “not found”

```

 $T \leftarrow S;$ 
while  $|T| > 1$  do
   $x \leftarrow \text{centralElement}(T);$ 
  if  $x$  is a maximal element of  $T$  then
     $G = \{y : y \text{ is a maximal element of } T\};$ 
  else
     $G = \{y : y \in \text{siblings}(x)\};$ 
  end if
   $y \leftarrow$  result of exact Grover search on  $G;$ 
  if result is “not found” then
     $T \leftarrow T \setminus \{z : \exists y' \in G, z \leq y'\};$ 
  else
     $T \leftarrow \{z : z \leq y\};$ 
  end if
end while
if  $|T|=1$  then
  return single element in  $T;$ 
else
  return not found;
end if

```

We will now prove an upper bound on the query complexity of Algorithm 1, via a couple of preparatory lemmas.

Lemma 3.4. *In each iteration of the loop, the total weight of the nodes in G is at least*

$|T|/2$.

Proof. If x is a maximal element, then the total weight of the nodes in G is clearly $|T|$, as every maximal element is added. If x is covered by an element p , then the total weight of the nodes in G will be $wt(p) - 1$. But $wt(p) > \lceil |T|/2 \rceil$ (as otherwise p would have been returned by `centralElement` rather than x), so we are done. \square

Lemma 3.5. *In each iteration of the loop, $|G| \leq 1/\gamma^S$.*

Proof. We will show that $\gamma^G = 1/|G|$, implying $\gamma^S \leq 1/|G|$. Restrict the marked element to being an element of G . Then an algorithm can only remove elements of G from consideration by querying within G . This is because, if x is not a maximal element of T , all the members of G are covered by a single element p , so the only queries that can allow us to reject members of G are queries to members of G . Alternatively, if x is a maximal element of T , then it is easy to see that x is actually also a maximal element of S . So G will contain all the maximal elements of S , and again the only queries that can allow us to reject members of G are queries to members of G . \square

Theorem 3.6. *Algorithm 1 finds the marked element in a forest-like n -element poset S , or outputs that no such element exists, with certainty using at most $O\left(\log n/\sqrt{\gamma^S}\right)$ queries to S .*

Proof. It is immediate that the algorithm is correct, as each iteration of the loop is guaranteed to remove at least one element from T . It remains to prove an upper bound on its query complexity. If the marked element a is in the set T at all, we are guaranteed that either $a \leq x$ for either exactly one element $x \in G$, or for no elements in G . The Grover search step will thus either reduce the search space to the elements $\{z\}$ of T for which $z \leq x$, or will remove all the elements $z \in T$ that are less than any element in G from consideration. Each element of G has weight at most $\lceil |T|/2 \rceil$, and by Lemma 3.4, their total weight is at least $|T|/2$. So each iteration of the loop will reduce the size of T by at least about half. By Lemma 3.5, each Grover search uses at most $O(1/\sqrt{\gamma^S})$ queries, so the theorem is proven. \square

In some cases, Algorithm 1 may do better than this upper bound suggests. One such example is searching a completely unstructured set (in which case the algorithm reduces to standard unstructured search, and thus achieves an $O(\sqrt{n}) = O\left(1/\sqrt{\gamma^S}\right)$ query complexity). As another example, it is easy to convince oneself that Algorithm 1 finds the marked element in a poset whose Hasse diagram is a complete k -ary tree with l levels using $O(\sqrt{kl})$ queries, rather than the $O(\sqrt{kl} \log k)$ queries guaranteed by Theorem 3.6.

Finally, note that the extension to searching for an unknown number of marked elements is straightforward: in this case, the exact Grover search step is replaced by picking an element y from G uniformly at random. If there exists a marked element a such that $a \leq y'$ for some element $y' \in G$, then the probability that $y = y'$ is at least $1/\sqrt{\gamma^S}$. We need to boost this success probability to $\Omega(1 - 1/\log n)$ in order for the success probability

after $O(\log n)$ recursions to be $\Omega(1)$. By a result of Buhrman et al [11] on amplification of classical probabilistic algorithms with one-sided error, this can be achieved using $O(\sqrt{\log \log n}/\sqrt{\gamma^S})$ iterations of picking $y \in G$ uniformly at random, giving an overall complexity of $O\left(\log n \sqrt{\log \log n}/\sqrt{\gamma^S}\right)$.

4 The concrete model

In this section, we consider the problem of poset search in the model studied by Linial and Saks [20], where the poset is thought of as storing partially sorted integers (or elements from any other totally ordered set), and querying an element of the poset returns the integer stored at that element. Note that we redefine $D(S)$, $Q_E(S)$ and $Q_2(S)$ appropriately.

4.1 Overall relationships

This model appears more complex to analyse, as the complexity of the search problem now depends not only on the structure of the poset being searched, but also on the integers that are stored in that poset. Also, the classical analysis of Linial and Saks [20] relies on certain properties of classical algorithms for poset search that quantum algorithms seem not to share. For example, at the end of a correct classical algorithm which searched unsuccessfully for the element a in S , every element $x \in S$ must have been classified according to whether $x < a$, $x = a$ or $x > a$. Quantum algorithms appear not to have this property.


However, we can develop a quantum lower bound that is similar to a known classical lower bound based on the size of the largest “unsorted” subset of S , namely the size of the largest antichain, $w(S)$. It turns out that finding an element in such a subset reduces to an unstructured search problem. We begin with a lemma whose classical part was shown by Linial and Saks [20] with a different proof.

Lemma 4.1. *Let S be a poset and let T be a section of S . Then $D(S) \geq D(T)$, $Q_E(S) \geq Q_E(T)$ and $Q_2(S) \geq Q_2(T)$.*

Proof. First, note that S can be partitioned into three disjoint subsets (or *layers*): the set T ; an “upper” set U where for all $u \in U$, there is no $t \in T$ such that $u \leq t$; and a “lower” set V where for all $v \in V$, there is no $t \in T \cup U$ such that $t \leq v$. Assume S has n elements, identified with the integers. Let V store the integers $\{1, \dots, |V|\}$ in some manner consistent with its partial order, and similarly let U store the integers $\{|V| + |T| + 1, \dots, n\}$. By the definition of the partitioning of S , T can store every permutation of the integers $\{|V| + 1, \dots, |V| + |T|\}$ that is consistent with its own partial order, independently of the integers stored in the remainder of S .

Now consider an adversarial strategy where the marked element is guaranteed to be in the set $\{|V| + 1, \dots, |V| + |T|\}$, and thus is stored in T . Any query to elements in U or V will then give no information about the position of the marked element within T , so any classical or quantum algorithm can restrict itself to making queries to elements in T . But any classical [exact quantum, bounded-error quantum] algorithm to find a marked element in T that only makes queries to elements in T must use $D(T)$ [$Q_E(T)$, $Q_2(T)$] queries. \square

Note that this property does not hold for arbitrary subsets of posets [20]: for example, the following posets $S, T \subset S$ have $D(S) = 3$ but $D(T) = 4$. The theorem does not hold at all in the abstract model of poset search discussed in the previous section.



Lemma 4.2. *Let S be an n -element unstructured poset. Then $D(S) = n$ and $Q_2(S) = \Omega(\sqrt{n})$.*

Proof. Let S store an arbitrary permutation π of the integers $\{1, \dots, n\}$, and let the marked element be $a = \pi(1)$. The classical lower bound is obvious [20] (as the only information obtained from a query to an element $x \in S$ is whether $a = x$ or $a \neq x$, every element in S may need to be queried in the worst case). In the quantum case, the lower bound of Ambainis on inverting a permutation [3] may be used to show that any quantum algorithm to find a requires $\Omega(\sqrt{n})$ queries. \square

Theorem 4.3. *Let S be an n -element poset. Then $D(S) = \Omega(w(S))$ and $Q_2(S) = \Omega(\sqrt{w(S)})$. Also, $Q_2(S) = \Omega(\log n)$.*

Proof. Let T be the largest antichain in S . T is unstructured, T is a section of S and $|T| = w(S)$. The first part of the theorem follows immediately from Lemma 4.1 and Lemma 4.2. For the second part, note that any quantum algorithm to find a marked element in S could also be used to find a marked element in a totally ordered set of n elements. The lower bound then follows from the lower bound of Ambainis [5] (improved by Høyer, Neerbek, and Shi [18]) on quantum search of an ordered list. \square

We now consider the question of upper bounds. It turns out that, up to a logarithmic factor, the width $w(S)$ *completely* characterises the classical and quantum query complexities of searching in this model. To show this, we will need the following powerful combinatorial result, which says something about the decomposition of a poset into chains.

Theorem 4.4. (Dilworth’s Theorem [15])

Let S be an n -element poset with $w(S) = k$. Then S is the union of k disjoint chains.

In fact, such a decomposition can be found in time $O(n^3)$ [8].

Lemma 4.5. *Let S be a poset. Then we have $D(S) = O(w(S) \log h(S))$ and $Q_E(S) = O(\sqrt{w(S)} \log h(S))$.*

Proof. Decompose S into a set C containing $w(S)$ disjoint chains, each of which contains at most $h(S)$ elements. The classical algorithm proceeds by searching each chain in C in turn, using binary search. The total number of queries required is therefore $O(w(S) \log h(S))$.

In the quantum case, our algorithm will nest an exact binary search algorithm within the exact variant of Grover’s search algorithm. We produce an oracle P_a which, when given

the identifier of a chain in C as input, returns whether the desired element a is contained within that chain; each call to P_a clearly requires at most $O(\log h(S))$ queries to the set. As the chains are disjoint, we are guaranteed that P_a will return 1 on only one input. The exact variant of Grover’s algorithm therefore requires (see Theorem 2.1) $O(\sqrt{w(S)})$ queries to P_a to determine which chain (if any) contains a . A final $O(\log h(S))$ queries are used to find a within that chain, for an overall query complexity of $O(\sqrt{w(S)} \log h(S))$. \square

If the binary search parts of this algorithm are replaced by the use of a quantum ordered search algorithm (e.g. [14]), the query complexity can be improved by a constant factor. Note that this algorithm actually also works in the abstract model of poset search, thus showing that, as one might expect, search in the abstract model is always at least as easy as in the concrete model (up to the $\log h(S)$ factor). Furthermore, note that an extension to search where a given integer may be stored at multiple positions in the poset is immediate: the Grover search steps are replaced by search for an unknown number of marked items [9] to give an $O(\sqrt{w(S)} \log h(S))$ -query bounded-error quantum algorithm.

We can now show that the classical and quantum query complexities of poset search in the concrete model are polynomially related.

Theorem 4.6. *Let S be an n -element poset with $Q_2(S) = k$. Then $D(S) = O(k^2 \log n) = O(k^3)$.*

Proof. Follows immediately from the quantum lower bounds of Lemma 4.3 and the classical upper bound of Lemma 4.5. \square

4.2 Searching a partially sorted array

Consider the following problem. We are given a d -dimensional $m_1 \times m_2 \times \dots \times m_d$ array of integers T that has been sorted in ascending order in each dimension (i.e. $(i_1 \leq j_1) \wedge (i_2 \leq j_2) \wedge \dots \wedge (i_d \leq j_d) \Rightarrow T(i_1, \dots, i_d) \leq T(j_1, \dots, j_d)$), and must find a given integer in this array, or output “not found”, using the minimum number of queries to the array. It is easy to see that this structure gives rise to a partially ordered set; see Figure 2 for the Hasse diagram of such a poset.



Figure 2: A 3×3 2-dimensional array sorted by rows and columns, and its corresponding Hasse diagram.

We are particularly interested in the special case where $m_i = m$ for all i . Call the poset corresponding to such a d -dimensional array $S_{d,m}$. Linial and Saks give [20] an $O(m^{d-1})$ classical algorithm for the problem of searching $S_{d,m}$, which is asymptotically optimal.

When $d = 2$, it is easy to see that we have $w(S_{2,m}) = m$. For higher d , Linial and Saks show that $w(S_{d,m}) = \Theta(m^{d-1})$. This follows from consideration of the set of elements that are indexed by a position (i_1, \dots, i_d) such that $\sum_k i_k = m + 1$; this is clearly an antichain and can be shown to have size $\Theta(m^{d-1})$. It is thus immediate from Lemma 4.5 and Lemma 4.3 that there exists a quantum algorithm that searches this poset using $O(m^{(d-1)/2} d \log m)$ queries, which is optimal up to the $d \log m$ factor.

We can write down such an algorithm explicitly as follows. The algorithm for $d = 1$ is just binary search. For $d = 2$, we nest a binary search algorithm on the rows within Grover search on the columns for an overall query complexity of $O(\sqrt{m} \log m)$. For $d = 3$, the algorithm simply performs Grover search on m copies of the $d = 2$ search algorithm, giving $O(m \log m)$ queries, and so on for $d > 3$.

It is worth noting that this poset structure is an example where searching in the abstract model is significantly easier than in the concrete model. Indeed, there exists a simple $O(d \log m)$ classical algorithm for search in the abstract model: simply perform binary search on each dimension of T .

In the following section, we will give an asymptotically optimal bounded-error quantum algorithm that searches a 2-dimensional $m \times m$ array of *distinct* integers in $O(\sqrt{m})$ queries. This then implies an asymptotically optimal $O(m^{(d-1)/2})$ -query algorithm for searching a d -dimensional $m \times m \times \dots \times m$ array of distinct integers. The optimal d -dimensional algorithm follows from treating the array as the union of m^{d-2} disjoint 2-dimensional $m \times m$ arrays. Each 2-dimensional array is searched by the optimal algorithm, which is treated as an oracle within an overall application of quantum search. Although the 2-dimensional search algorithm is bounded-error, a version of quantum search which can cope with bounded-error inputs (due to Høyer, Mosca and de Wolf [19]) can be used to achieve a constant probability of success in $O(m^{(d-1)/2})$ queries.

4.2.1 Optimal search of a 2-dimensional array sorted by rows and columns

In this section, we give an asymptotically optimal algorithm to search for a known integer a within an $r \times c$ 2-dimensional array of distinct integers sorted by rows and columns. We will start by describing a classical algorithm for the same problem, which is asymptotically (but not exactly [20]) optimal. The algorithm's operation will be described in terms of the original array, rather than the more abstract poset structure. Call the $\lceil \frac{r}{2} \rceil$ 'th row of the array the *central* row R , and similarly let the $\lceil \frac{c}{2} \rceil$ 'th column be the central column C .

If $r \leq c$, begin by performing binary search for a on the central column, using $O(\log r)$ queries. If $r > c$, do the same, but on the central row, using $O(\log c)$ queries. Assume $r \leq c$ and that a is not in the central column (otherwise, a will be found by the binary search, and can be returned immediately). Then by the end of the binary search we will have found an element x such that $x = \max_{x' \in C} (x' < a)$, and an element y such that $y = \min_{y' \in C} (y' > a)$ (so y is positioned directly below x in the array). This then implies that all elements in the array above and to the left of x are also less than a , and similarly all elements below and to the right of y are greater than a , so these elements can be discarded. As x and y are in the central column, we must have excluded at least half of the elements in the array from

consideration.

We are then left with two smaller instances of the same problem to solve: the subarray below and to the left of y , and the subarray above and to the right of x . The algorithm proceeds to search these subarrays recursively until a is found, performing binary search on central rows or central columns as appropriate.

1	3	5	10	13
2	4	7	11	14
6	8	9	15	21
12	16	17	20	24
18	19	22	23	25

1	3	5	10	13
2	4	7	11	14
6	8	9	15	21
12	16	17	20	24
18	19	22	23	25

1	3	5	10	13
2	4	7	11	14
6	8	9	15	21
12	16	17	20	24
18	19	22	23	25

Figure 3: Example of the classical algorithm’s operation when searching for the element 11: dark grey squares are those that are searched in each round, light grey squares have been excluded from consideration, white squares are still to be searched. Here, 11 is found with only 2 levels of recursion.

How many queries to the array does this algorithm require? Let $T(m)$ denote the number of queries used to search an $r \times c$ array, with $m = \max(r, c)$. Then it is easy to see that $T(m)$ will be maximised if each level of binary search always terminates as close to the centre of the central column/row as possible (thus maximising the number of queries required for binary search in the next level of recursion). We therefore have

$$T(m) \leq \lceil \log_2 m + 1 \rceil + 2T(m/2) \tag{5}$$

and unwinding the recursion gives $T(m) = O(m)$.

We would like to find an analogous quantum algorithm that achieves some reduction in queries by searching the subarrays in superposition, rather than sequentially. In fact, it turns out that we can make a general statement about when recursive classical search algorithms can be turned into improved quantum search algorithms, which is given as the following lemma. The proof is a fairly straightforward generalisation of a powerful result of Aaronson and Ambainis [1], so is deferred to Appendix A.

Lemma 4.7. *Let P_n be the problem of searching an abstract database, parametrised by an abstract size n , for a known element which may or may not be in the database. Let $T(n)$ be the time required for a bounded-error quantum algorithm to solve P_n , i.e. to find the element, or output “not found”. Let P_n satisfy the following conditions:*

- *If $n \leq n_0$ for some constant n_0 , then there exists an algorithm to find the element, if it is contained in the database, in time $T(n) \leq t_0$, for some constant t_0 .*
- *If $n > n_0$, then the database can be divided into k sub-databases of size at most $\lceil n/k \rceil$, for some constant $k > 1$.*
- *If the element is contained in the original database, then it is contained in exactly one of these sub-databases.*

- Each division into sub-databases uses time $f(n)$, where $f(n) = O(n^{1/2-\epsilon})$ for some $\epsilon > 0$.

Then $T(n) = O(\sqrt{n})$.

We show that the search problem in question fits the conditions of the lemma. We consider the problem to be parametrised by a “size” $m = \max(r, c)$. Assuming that a is stored in the set and is not stored in the central row/column, one step of the classical procedure given above will divide any array of size m into two arrays of size at most $\lceil m/2 \rceil$, exactly one of which contains a , in time $O(\log m)$. This division can be performed recursively until the arrays are reduced to a constant size. In the case where the binary search of the central row/column actually finds a , the algorithm can easily be modified to not return a immediately, but to restrict the search area in the next recursion to two subarrays, exactly one of which includes a , and both of which are of size at most $\lceil m/2 \rceil$.

There thus exists a quantum algorithm, given explicitly in Appendix A, that can find an arbitrary element a in the array in $O(\sqrt{m})$ time, and hence $O(\sqrt{m})$ queries.

4.2.2 Finding the intersection of two increasing lists

Classically, there is a correspondence between the problem of searching in an $r \times c$ array sorted by rows and columns and merging two sorted lists of length r and c : any decision tree for the one problem gives a decision tree for the other [20]. However, this does not appear to hold for quantum algorithms; indeed, it is straightforward to show, using Holevo’s Theorem [17], an $\Omega(r + c)$ quantum query lower bound for the merge problem. Nevertheless, we can define a natural search problem that turns out to arise from the poset search problem.

Problem: Given two lists of integers in increasing order, output an integer that occurs in both lists, or report that no such integer exists.

This can be thought of as a special case of the element distinctness problem [2]. It was studied by Buhrman et al [10], who also refer to it as the *monotone claw* problem (a claw is an input on which two functions take the same value). Let the lists be denoted L and M and be of length l and m respectively, with $l \geq m$. Then the ingenious algorithm of [10] finds an integer occurring in both lists using $O(\sqrt{l}c^{\log^* l})$ queries, where \log^* is the iterated logarithm function and c is a constant. This algorithm can easily be translated into the setting of poset search, and allows an $m \times m$ array that is sorted by rows and columns, and may contain duplicates, to be searched using $O(\sqrt{m}c^{\log^* m})$ time for some constant c .

Here, we will go in the other direction, and show that the algorithm of Section 4.2.1 can be used to find the integer occurring in both sorted lists using $O(\sqrt{l})$ time. As noted in [10], there is an $\Omega(\sqrt{l})$ lower bound for this problem, so the algorithm given here is asymptotically optimal. However, as $c^{\log^* l}$ is already a near-constant function, the new algorithm may be only of theoretical interest, and we describe it briefly.

Consider a notional $l \times m$ array T where entry $T(x, y)$ contains the value $L_x - M_{m+1-y}$. Querying one entry of T uses one query to each list. As the entries in L and M are in increasing order, it is easy to see that T is increasing along rows and columns, and that finding a 0 entry in T corresponds to finding an element of L that also occurs in M . Call

such an element a *match*. If there is only one match, it is immediate that the algorithm of the previous section can be used to find the single 0 entry in T , or output that no such entry exists, in time $O(\sqrt{l})$.

There are two possible reasons for there being more than one match. The first is that L and M may contain duplicate elements (i.e. may be increasing but not strictly increasing). If this is the case, and if one of the duplicate elements in L (say) is also in M , there will be a contiguous rectangle of 0 entries in the array T (call this a *zero block*), rather than a single 0. Assume that there is only one zero block. Then the algorithm of Section 4.2.1 must be modified to ensure that, after any splitting of the array into two subarrays, at most one of these arrays contains a 0 entry; i.e. to ensure that the zero block does not get split across subarrays. This is necessary to ensure that the conditions of Lemma 4.7 are satisfied. It is easy to see that, in each round of recursion, the zero block can only be split if it lies across a row or column that is used for binary search in that recursion. In order to ensure that only one of the two subarrays produced contains part of the zero block in this case, the binary search of a row (column) can simply be modified to split on the first or last zero entry in that row (column), with no change to the asymptotic complexity. Call this new algorithm the single-block algorithm.

The second case where there may be more than one match is when there is more than one element in L that also occurs in M (or vice versa). In this case, the idea (inspired by [1]) is to reduce the problem to searching for a single zero block by probabilistically removing elements from the lists. The extended algorithm first runs the single-block algorithm. Assuming that this algorithm outputs “not found”, the next step is to produce a new pair of smaller lists $L^{(2)}$ and $M^{(2)}$, which will give rise to a notional array $T^{(2)}$, where $T^{(2)}(x, y) = L_x^{(2)} - M_{m+1-y}^{(2)}$.

The reduction in size is achieved by first splitting each list into chunks of size 2. One element (picked at random) within each chunk of L is included in $L^{(2)}$, and similarly for M and $M^{(2)}$. The single-block algorithm is then run on these smaller lists. Assuming that the result is again “not found”, the chunk size is doubled to 4, and the process repeats, using a chunk size of 2^k in each round k . Assuming that the single-block algorithm does not find a match in any of the $O(\log l)$ rounds, the final output is “not found”. The time required for this overall algorithm is then bounded by $O\left(\sum_k \sqrt{l/2^k}\right) = O(\sqrt{l})$.

We sketch a proof that this algorithm succeeds with constant probability. First, it is easy to see that there can be at most one zero block in each row and column of the array $T^{(k)}$ in any round k . Using this, one can show that, if there are z zero blocks in T , the probability that exactly one remains in $T^{(k)}$ is at least $z/2^{2k}(1 - z/2^{2k})$. If we take $k = \lceil \log z/2 \rceil + 1$, this is lower bounded by a constant, so for any z the single-block algorithm succeeds with constant probability in at least one round.

5 Conclusions

We have given general upper and lower bounds on quantum search of partially ordered sets, in two different models. Satisfyingly, in the two cases where results were already known

on poset search (i.e. totally ordered sets and unstructured sets), our lower bounds reduce to known lower bounds, and our new quantum algorithms are (asymptotically) as efficient as the known most efficient algorithms. The bounds in the concrete model are perhaps particularly interesting, because they follow from decomposing a poset into “structured” and “unstructured” components, and show that, intuitively, almost all the speed-up that can be obtained from quantum search of a poset S is obtained from searching the unstructured parts of S .

Although we concentrated on the model of query complexity, our quantum algorithms in both models are efficient in the sense that, given a poset S to be searched, quantum circuits for the algorithms given here can be produced in time polynomial in the size of S . Also, the non-query transformations used by the algorithms given here are efficiently implementable.

However, there are still several open questions. Firstly: in the abstract model, is there a general lower bound of $Q(S) = \Omega(\log n)$? This would be an interesting generalisation of the known logarithmic quantum lower bound on searching an ordered list [5, 18]. Also, can the logarithmic factors in the quantum upper bounds in both models be improved, perhaps by being changed into additive terms?

There are several possible extensions involving search for multiple marked elements. In the abstract model, can a $O\left(\log n / \sqrt{\gamma^S}\right)$ -query algorithm be produced for search for multiple marked elements in arbitrary posets? In the concrete model, could the algorithm of section 4.2.1 be extended to arrays that may contain duplicate elements?

Acknowledgements

I would like to thank Richard Jozsa, Raphaël Clifford, Richard Low, Dan Shepherd and Aram Harrow for helpful discussions.

A Amplitude amplification of recursive search

The aim of this appendix is to give a proof of a somewhat generalised version of a powerful result that was shown by Aaronson and Ambainis [1] in the course of their work on quantum search of spatial regions. Informally, we would like to be able to find “cookbook” quantum algorithms for search problems for which there exists a recursive classical algorithm. We imagine that we are searching for a distinguished element in an abstract “database” that is parametrised by an abstract “size” n , which is some function of the number of elements in the database. We also imagine that we have the ability to search the database recursively: that is, in time given by some function $f(n)$, we can reduce the search problem to searching k instances of databases of size $\leq \lceil n/k \rceil$, for some constant $k > 1$.

It is straightforward to show that, classically, the marked element can be found deterministically in $O(n)$ time, by repeated use of this recursive search. An alternative probabilistic classical algorithm for this problem would be: split the input into a number of parts, pick one part uniformly at random, and call yourself recursively on that part. Our quantum

algorithm will apply amplitude amplification to this probabilistic algorithm. It will turn out to be advantageous to only amplify a small number of times within the recursive algorithm, and then to amplify again at the end. Amplifying to high probabilities too soon is less efficient [1]; conversely, if amplitude amplification were only applied at the end of the algorithm, we would require $\Omega(\sqrt{n})$ iterations to amplify the probability to a constant. If the process of dividing the input required time $f(n) = \omega(1)$, this would hurt the overall complexity.

The fundamental amplitude amplification result of Brassard et al [9] states that, given a quantum algorithm A with success probability ϵ , we can achieve a success probability of $\Omega(1)$ with only $O(1/\sqrt{\epsilon})$ uses of A . However, here we will need a tighter analysis due to Aaronson and Ambainis [1], as constants are important within the recursive algorithm.

Lemma A.1. *Given a quantum algorithm with success probability at least ϵ , then by executing it $t = 2m + 1$ times, where $m \leq \pi/(\arcsin \sqrt{\epsilon}) - 1/2$, we can achieve success probability at least $(1 - \frac{1}{3}t^2\epsilon)t^2\epsilon$.*

We are now ready to give a formal definition of a quantum algorithm for recursive search problems, and to upper-bound its time complexity. The algorithm and its analysis closely follow the results on spatial search of a d -dimensional cube of [1].

Lemma A.2. *Let P_n be the problem of searching an abstract database, parametrised by an abstract size n , for a known element which may or may not be in the database. Let $T(n)$ be the time required for a bounded-error quantum algorithm to solve P_n , i.e. to find the element, or output “not found”. Let P_n satisfy the following conditions:*

- *If $n \leq n_0$ for some constant n_0 , then there exists an algorithm to find the element, if it is contained in the database, in time $T(n) \leq t_0$, for some constant t_0 .*
- *If $n > n_0$, then the database can be divided into k sub-databases of size at most $\lceil n/k \rceil$, for some constant $k > 1$.*
- *If the element is contained in the original database, then it is contained in exactly one of these sub-databases.*
- *Each division into sub-databases uses time $f(n)$, where $f(n) = O(n^{1/2-\epsilon})$ for some $\epsilon > 0$.*

Then $T(n) = O(\sqrt{n})$.

Proof. Our quantum algorithm will be parametrised by two constants α and δ , whose values we will take to be $\delta = \epsilon/2$, $\alpha = \frac{\epsilon(4-3\epsilon)}{8(2-\epsilon)}$, and will be based on the following probabilistic classical algorithm:

If $n \leq n_0$, then find the desired element directly or output “not found” (using at most t_0 steps). Otherwise, assume that there exists an integer l such that $n^\delta = k^l$ ². Recursively

²We assume here that l and n^α are integers. One can show that the need to round these quantities up or down has no effect on the overall asymptotic complexity.

divide the problem into subproblems l times, leaving n^δ subproblems, each of size at most $n^{1-\delta}$. Pick one of the parts at random, and call yourself recursively on that part. Repeat until the desired element has been found.

We will perform a number of iterations of amplitude amplification on this algorithm such that it is executed n^α times. Then we have

$$T(n) \leq n^\alpha \left(\sum_{i=0}^{l-1} k^i f(n/k^i) + T(n^{1-\delta}) \right) \quad (6)$$

$$\leq n^\alpha \left(\ln^\delta f(n) + T(n^{1-\delta}) \right) \quad (7)$$

$$= n^\alpha f'(n) + n^{\alpha(1+(1-\delta))} f'(n^{1-\delta}) + n^{\alpha(1+(1-\delta)+(1-\delta)^2)} f'(n^{(1-\delta)^2}) + \dots + t_0 \quad (8)$$

$$= O(n^{\alpha(1+(1-\delta)+(1-\delta)^2+\dots)}) \quad (9)$$

$$= O(n^{\alpha/\delta}) \quad (10)$$

where we define $f'(n) = \ln^\delta f(n) = O(n^{(1-\epsilon)/2} \log n)$. The fourth line follows because $(1-\epsilon)/2 < \alpha(1/\delta - 1)$, so for any $m \geq 0$ we have $f'(n^{(1-\delta)^m}) = O(n^{(1-\delta)^m(1-\epsilon)/2} \log n) = o(n^{(\alpha/\delta)(1-\delta)^{m+1}})$, so the $f'(n^{(1-\delta)^m})$ parts of the third line are negligible.

We now calculate a lower bound on the probability of success $P(n)$ of this algorithm. If there were no amplification, we would have $P(n) \geq n^{-\delta} P(n^{1-\delta})$ for $n > n_0$, and $P(n) = 1$ for $n \leq n_0$. So, by Lemma A.1, we have

$$P(n) \geq (1 - n^{2\alpha-\delta}/3) n^{2\alpha-\delta} P(n^{1-\delta}) \quad (11)$$

$$= [(1 - n^{2\alpha-\delta}/3)(1 - n^{(2\alpha-\delta)(1-\delta)}/3) \dots] n^{(2\alpha-\delta)(1+(1-\delta)+(1-\delta)^2+\dots)} \quad (12)$$

$$= [(1 - n^{2\alpha-\delta}/3)(1 - n^{(2\alpha-\delta)(1-\delta)}/3) \dots] \Omega(n^{2\alpha/\delta-1}) \quad (13)$$

We claim that the remaining product of bracketed terms is lower bounded by a constant that does not depend on n . First, note that the algorithm recurses R times, for some $R = O(\log \log n)$. Now

$$\prod_{k=0}^R \left(1 - \frac{1}{3} n^{(2\alpha-\delta)(1-\delta)^k}\right) \geq 1 - \frac{1}{3} \sum_{k=0}^{O(\log \log n)} n^{(2\alpha-\delta)(1-\delta)^k} \geq 1 - O(n^{2\alpha-\delta} \log \log n) = 1 - o(1) \quad (14)$$

giving the result $P(n) = \Omega(n^{2\alpha/\delta-1})$.

By wrapping this algorithm in another level of amplitude amplification, we can use $O(P(n)^{-1/2})$ iterations of it to achieve a constant probability of success of finding the marked element in time $O(T(n)P(n)^{-1/2}) = O(n^{\alpha/\delta} n^{1/2-\alpha/\delta}) = O(\sqrt{n})$. \square

References

- [1] S. Aaronson, A. Ambainis. Quantum search of spatial regions. *Theory of Computing* 1, pp. 47-79, [quant-ph/0303041](#), 2005.
- [2] S. Aaronson, Y. Shi. Quantum lower bounds for the collision and the element distinctness problems. *Journal of the ACM* 51, pp. 595-605, 2004.

- [3] A. Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences* 64, pp. 750-767, [quant-ph/0002066](#), 2002.
- [4] A. Ambainis, K. Iwama, A. Kawachi, H. Masuda, R. Putra, S. Yamashita. Quantum identification of Boolean oracles. *Proc. STACS'04*, pp. 93-104, [quant-ph/0403056](#), 2004.
- [5] A. Ambainis. A better lower bound for quantum algorithms searching an ordered list. *Proc. FOCS'99*, pp. 352-357, [quant-ph/9902053](#), 1999.
- [6] A. Atici, R. Servedio. Improved bounds on quantum learning algorithms. *Quantum Information Processing* 4, pp. 355-386, [quant-ph/0411140](#), 2005.
- [7] Y. Ben-Asher, E. Farchi, I. Newman. Optimal search in trees. *SIAM J. Comput.* 28, pp. 2090-2102, ECCC TR96-044, 1999.
- [8] K. Bogart. *Introductory combinatorics* (3rd edition). Brooks Cole, 2000.
- [9] G. Brassard, P. Høyer, M. Mosca, A. Tapp. Quantum amplitude amplification and estimation. *Quantum Computation and Quantum Information: A Millennium Volume*, AMS Contemporary Mathematics Series, [quant-ph/0005055](#), 2002.
- [10] H. Buhrman, C. Durr, M. Heiligman, P. Høyer, F. Magniez, M. Santha, R. de Wolf. Quantum algorithms for element distinctness. *SIAM J. Comput.* 34, pp. 1324-1330, [quant-ph/0007016](#), 2005.
- [11] H. Buhrman, R. Cleve, R. de Wolf, C. Zalka. Bounds for small-error and zero-error quantum algorithms. *Proc. FOCS'99*, pp. 358-368, [cs/9904019](#), 1999.
- [12] R. Carmo, J. Donadelli, Y. Kohaykawa, E. Laber. Searching in random partially ordered sets. *Proc. LATIN'02*, pp. 278-292, 2002.
- [13] N. Cerf, L. Grover, C. Williams. Nested quantum search and structured problems, *Phys. Rev. A* 61 032303, [quant-ph/9806078](#), 2000.
- [14] A. Childs, A. Landahl, P. Parrilo. Improved quantum algorithms for the ordered search problem via semidefinite programming. [quant-ph/0608161](#), 2006.
- [15] R. P. Dilworth. A decomposition theorem for partially ordered sets. *The Annals of Mathematics* 51, pp. 161-166, 1950.
- [16] L. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.* 79, pp. 325-328, [quant-ph/9706033](#), 1997.
- [17] A. S. Holevo. Bounds for the quantity of information transmittable by a quantum communications channel. *Problemy Peredachi Informatsii*, vol. 9, no. 3, pp. 3-11, 1973. English translation *Problems of Information Transmission*, vol. 9, pp. 177-183, 1973.
- [18] P. Hyer, J. Neerbek, Y. Shi. Quantum complexities of ordered searching, sorting, and element distinctness. *Proc. ICALP'01*, pp. 62-73, [quant-ph/0102078](#), 2001.

- [19] P. Hyer, M. Mosca, R. de Wolf. Quantum search on bounded-error inputs. *Proc.ICALP'03*, pp. 291-299, [quant-ph/0304052](#), 2003.
- [20] N. Linial, M. Saks. Searching ordered structures. *Journal of Algorithms* 6, pp. 86-103, 1985.
- [21] N. Linial, M. Saks. Every poset has a central element. *J. Comb. Th. Ser. A* 40, pp. 195-210, 1985.
- [22] G. Long. Grover algorithm with zero theoretical failure rate. *Phys. Rev. A* 64 022307, [quant-ph/0106071](#), 2001.
- [23] K. Onak, P. Parys. Generalization of binary search: searching in trees and forest-like partial orders. *Proc. FOCS'06*, 2006.
- [24] R. A. Servedio, S. J. Gortler. Quantum versus classical learnability. *Proc. CCC'01*, pp. 138-148, [quant-ph/0007036](#), 2001.
- [25] C. Zalka. Grover's quantum searching algorithm is optimal. *Phys. Rev. A* 60, pp. 2746-2751, [quant-ph/9711070](#), 1999.