

Lower bounds for predecessor searching in the cell probe model [☆]

Pranab Sen ^{a,*,1,2}, S. Venkatesh ^{b,3}

^a *NEC Laboratories America, 4 Independence Way, Suite 200, Princeton, NJ 08540, USA*

^b *Department of Computer Science, University of Victoria, Victoria, BC V8W 3P6, Canada*

Received 13 November 2003; received in revised form 28 June 2006

Available online 13 June 2007

Abstract

We consider a fundamental problem in data structures, *static predecessor searching*: Given a subset S of size n from the universe $[m]$, store S so that queries of the form “What is the predecessor of x in S ?” can be answered efficiently. We study this problem in the cell probe model introduced by Yao [A.C.-C. Yao, Should tables be sorted, *J. Assoc. Comput. Mach.* 28 (3) (1981) 615–628]. Recently, Beame and Fich [P. Beame, F. Fich, Optimal bounds for the predecessor problem and related problems, *J. Comput. System Sci.* 65 (1) (2002) 38–72] obtained optimal bounds as functions of either m or n only on the number of probes needed by any deterministic query scheme if the associated storage scheme uses only $n^{O(1)}$ cells of word size $(\log m)^{O(1)}$ bits.

We give a new lower bound proof for this problem that matches the bounds of Beame and Fich. Our lower bound proof has the following advantages: it works for randomised query schemes too, while Beame and Fich’s proof works for deterministic query schemes only. In addition, it is simpler than Beame and Fich’s proof. In fact, our lower bound for predecessor searching extends to the ‘quantum address-only’ query schemes that we define in this paper. In these query schemes, quantum parallelism is allowed only over the ‘address lines’ of the queries. These query schemes subsume classical randomised query schemes, and include many quantum query algorithms like Grover’s algorithm [L. Grover, A fast quantum mechanical algorithm for database search, in: *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, 1996, pp. 212–219].

We prove our lower bound using the round elimination approach of Miltersen, Nisan, Safra and Wigderson [P. Bro Miltersen, Noam Nisan, S. Safra, A. Wigderson, On data structures and asymmetric communication complexity, *J. Comput. System Sci.* 57 (1) (1998) 37–49]. Using tools from information theory, we prove a strong round elimination lemma for communication complexity that enables us to obtain a tight lower bound for the predecessor problem. Our strong round elimination lemma also extends to quantum communication complexity. We also use our round elimination lemma to obtain a rounds versus communication tradeoff for the ‘greater-than’ problem, improving on the tradeoff in [P. Bro Miltersen, Noam Nisan, S. Safra, A. Wigderson, On data

[☆] This paper is the journal version of the conference papers [P. Sen, S. Venkatesh, Lower bounds in the quantum cell probe model, in: *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, in: *Lecture Notes in Comput. Sci.*, vol. 2076, Springer-Verlag, 2001, pp. 358–369] and [P. Sen, Lower bounds for predecessor searching in the cell probe model, in: *Proceedings of the 18th Annual IEEE Conference on Computational Complexity*, 2003, pp. 73–83].

* Corresponding author.

E-mail addresses: pgdsen@tcs.tifr.res.in (P. Sen), venkat@cs.uvic.ca (S. Venkatesh).

¹ Work done while the author was a graduate student at TIFR, Mumbai, India visiting UC Berkeley and DIMACS under a Sarojini Damodaran International Fellowship grant, and a postdoctoral researcher at LRI, Orsay, France supported by the EU 5th framework program QAIP IST-1999-11234 and by CNRS/STIC 01N80/0502 and 01N80/0607 grants.

² Present address: School of Technology and Computer Science, Tata Institute of Fundamental Research, Mumbai 400005, India.

³ Work done while the author was a postdoctoral researcher at IAS, Princeton, USA supported by NSF grant CCR-9987845 and a joint IAS-DIMACS post doctoral fellowship.

structures and asymmetric communication complexity, J. Comput. System Sci. 57 (1) (1998) 37–49]. We believe that our round elimination lemma is of independent interest and should have other applications.

© 2007 Elsevier Inc. All rights reserved.

Keywords: Predecessor searching problem; Cell probe model; Data structures; Lower bounds

1. Introduction

1.1. The problem and the model

A static data structure problem consists of a set of data D , a set of queries Q , a set of answers A , and a function $f : D \times Q \rightarrow A$. The aim is to store the data efficiently and succinctly, so that any query can be answered with only a few probes to the data structure. *Static predecessor searching* is a well studied problem in data structure design (see e.g. [2–4,28,29,38–40]). Data structures for answering predecessor queries can be used to construct data structures to answer other queries like rank (finding the number of elements in S that are smaller than or equal to x) and nearest neighbour (finding an element in S closest to x) efficiently. This motivates the need to design efficient data structures that support predecessor queries.

Let $[m]$ denote the set of integers $\{0, \dots, m - 1\}$.

Definition 1 (*Static predecessor searching*). In the problem of (m, n) -static predecessor searching, we are given a subset S of size n from the universe $[m]$. Our goal is to store the set S succinctly so that queries of the form “What is the predecessor of x in S ?” for $x \in [m]$ can be answered with a few probes to the data structure. If x has no predecessor in S , that is, x is smaller than every element in S , then return a default value, say, -1 .

In this paper, we study the static predecessor searching problem in Yao’s *cell probe* model [44]. The cell probe model is a natural and general model for proving upper and lower bound results in data structures.

Definition 2 (*The cell probe model*). An (s, w, t) *cell probe scheme* for a static data structure problem $f : D \times Q \rightarrow A$ has two components: a *storage scheme* and a *query scheme*. The storage scheme stores the data $d \in D$ as a table $T[d]$ of s cells, each cell of word size w bits. The storage scheme is classical deterministic. Given a query $q \in Q$, the query scheme computes $f(d, q)$ by making at most t probes to $T[d]$, where each probe reads one cell at a time, and the probes can be adaptive. In a deterministic cell probe scheme the query scheme is classical deterministic, in a randomised cell probe scheme it is classical randomised, and in an address-only quantum cell probe scheme it is bounded error address-only quantum (defined in Section 2.1).

Since in the cell probe model we only charge a scheme for the number of probes made to memory cells and for the total number of cells of storage used, and all internal computation is for free, lower bounds proved in the cell probe model hold in all reasonable data structure models (e.g. the unit cost RAM with the same word size) and give us insight into the intrinsic difficulty of the problem.

The goal is to design cell probe schemes for (m, n) -static predecessor searching using small space, i.e. $s = n^{O(1)}$ and $w = O(\log m)$, and at the same time making a small number of probes t in the worst case.

1.2. Previous work

We start by describing the sequence of results that lead to the currently best known upper bounds for the (m, n) -static predecessor problem. For a long time, the best upper bound known for the predecessor problem was due to the data structures of van Emde Boas et al. [38,39], and the data structures of Fredman and Willard [12]. In their papers, van Emde Boas et al. [38,39] gave a $(\Omega(m), O(\log m), O(\log \log m))$ deterministic cell probe solution for predecessor. The main drawback of their solution is that the number of cells used is very large. Later, Willard [40] reduced the number of cells used to $O(n)$. Building on the work of van Emde Boas et al. and Willard, Fredman and Willard [12], and Andersson [3] designed $(O(n), O(\log m), O(\sqrt{\log n}))$ deterministic cell probe schemes for

predecessor. Recently, Beame and Fich [4] improved on these upper bounds and showed a $(O(\frac{n^2 \log n}{\log \log n}), O(\log m), t)$ deterministic cell probe scheme for predecessor, where

$$t = \min \left\{ O\left(\frac{\log \log m}{\log \log \log m}\right), O\left(\sqrt{\frac{\log n}{\log \log n}}\right) \right\}.$$

The first lower bound for the (m, n) -static predecessor problem was proved by Ajtai [2], who showed that no $(n^{O(1)}, O(\log m), t)$ deterministic cell probe scheme for predecessor can have constant number of probes t . Miltersen [28] observed that there is a close connection between the cell probe complexity of a data structure problem and the communication complexity of a related communication game, and used this to improve Ajtai's lower bound to $\Omega(\sqrt{\log \log m})$ probes. Recently, building on Ajtai's and Miltersen's work, Beame and Fich [4] showed that their data structure described above is indeed optimal in the following sense: there is a constant c such that for all sufficiently large m a value of n can be chosen so that any $(n^{O(1)}, 2^{(\log m)^{1-\Omega(1)}}, t)$ deterministic cell probe scheme for predecessor must satisfy $t \geq \frac{c \log \log m}{\log \log \log m}$, as well as for all sufficiently large n a value of m can be chosen so that any $(n^{O(1)}, (\log m)^{O(1)}, t)$ deterministic cell probe scheme for predecessor must satisfy $t \geq \sqrt{\frac{c \log n}{\log \log n}}$. Similar lower bounds were proved by Xiao [41]. We would like to stress here that all the above lower bound proofs are complicated with many technical details. Also, they hold for deterministic cell probe schemes only.

The result of Beame and Fich gives rise to the following two questions: does their lower bound hold for randomised query schemes as well? It has been observed recently that randomisation enormously helps in the case of membership queries [5] and approximate nearest neighbour queries [6,16,25,26], and one might believe that it could help answer predecessor queries quickly as well. Secondly, is it possible to give a simple proof of the lower bound result of Beame and Fich?

A partial answer to both of the above questions was given by Miltersen, Nisan, Safra and Wigderson [29]. In their paper, they proved a general *round elimination* lemma for communication complexity. Using the connection between cell probe complexity of data structures and communication complexity, and their round elimination lemma, they showed the following lower bound for the predecessor problem: any $(n^{O(1)}, (\log m)^{O(1)}, t)$ randomised cell probe scheme for (m, n) -static predecessor must satisfy $t = \Omega(\sqrt{\log \log m})$ as a function of m , and $t = \Omega((\log n)^{1/3})$ as a function of n . Though the lower bound proved by [29] is weaker than that of [4], their approach had two advantages: their lower bound holds for randomised query schemes too, and the proof is much simpler. In their paper, Miltersen et al. ask if their round elimination based approach can be strengthened to obtain the lower bound of Beame and Fich.

1.3. Our results

We answer the question posed by Miltersen, Nisan, Safra and Wigderson. Our main result in this paper shows that the lower bound of Beame and Fich holds for address-only quantum cell probe schemes (and hence, for randomised cell probe schemes) as well.

1.3.1. The predecessor problem

Result 1. Suppose there is a $(n^{O(1)}, (\log m)^{O(1)}, t)$ randomised cell probe scheme for the (m, n) -static predecessor problem with error probability less than $1/3$. Then, there is a constant c such that:

- (a) For all sufficiently large m , a value of n can be chosen so that $t \geq \frac{c \log \log m}{\log \log \log m}$ as a function of m .
- (b) For all sufficiently large n , a value of m can be chosen so that $t \geq \sqrt{\frac{c \log n}{\log \log n}}$ as a function of n .

The same lower bound also holds for address-only quantum cell probe schemes for static predecessor searching.

We prove our lower bound for predecessor searching by combining the approach in [29] with a new round elimination lemma for communication complexity. Our round elimination lemma is a strengthening of the one proved in [29], and we believe it is of independent interest.

1.3.2. An improved round elimination lemma

In this paper, all communication protocols are two-party. The error probability of a randomised or quantum communication protocol is defined as the maximum error of the protocol for any input. For a general introduction to (classical) communication complexity, see the book by Kushilevitz and Nisan [24].

Let $f: \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ be any communication problem. Let us denote by $f^{(n),A}$ a new communication game in which Alice is given $x_1, x_2, \dots, x_n \in \mathcal{X}$, and Bob is given $y \in \mathcal{Y}$, $i \in [n]$ and also copies of x_1, x_2, \dots, x_{i-1} . Their task is to compute $f(x_i, y)$. $f^{(n),B}$ is defined similarly. Intuitively, if Alice starts the communication and her first message is much smaller than n bits, then she is unlikely to send much useful information about x_i to Bob as she is unaware of i . So it should be possible to eliminate the first message of Alice, giving rise to a protocol where Bob starts, with one less round of communication, and having similar message complexity and error probability. The round elimination lemma captures this intuition.

Definition 3. A $[t; l_1, l_2, \dots, l_t]^A$ ($[t; l_1, l_2, \dots, l_t]^B$) communication protocol is one where Alice (Bob) starts the communication, the i th message is l_i bits long, and the communication goes on for t rounds.

Result 2. Suppose the communication game $f^{(n),A}$ has a $[t; l_1, l_2, \dots, l_t]^A$ public coin randomised protocol with error less than δ . Then, f has a $[t-1; l_2, \dots, l_t]^B$ public coin randomised protocol with error less than $\epsilon \triangleq \delta + (1/2)(2l_1 \ln 2/n)^{1/2}$. A similar result holds for public coin quantum protocols (defined in Section 2.2) too.

The proof of this lemma uses tools from information theory. In particular, we use the *average encoding* theorem of Klauck, Nayak, Ta-Shma and Zuckerman [22]. Intuitively, this theorem says that if the mutual information between a random variable and its randomised encoding is small, then the probability distributions on code words for various values of the random variable are indeed close to the average probability distribution on code words.

1.3.3. Applications to other problems

We prove our lower bound result for predecessor by actually proving a lower bound for the *rank parity problem*. In the rank parity problem, we need to store a subset S of the universe $[m]$ so that given a query element $x \in [m]$, we can output whether the number of elements in S less than or equal to x is even or odd. Lower bounds for rank parity imply similar lower bounds for some other data structure problems like *point separation* [4] and *two-dimensional reporting range query* [29]. For details of the reduction from rank parity to the above problems, see the respective papers cited above.

Independently, the round elimination lemma has applications to problems in communication complexity. For example, let us consider communication protocols for the ‘greater-than’ problem GT_n in which Alice and Bob are given bit strings x and y respectively of length n each, and the goal is to find out if $x > y$ or not (treating x, y as integers between 0 and $2^n - 1$). Miltersen, Nisan, Safra and Wigderson [29], and Smirnov [37] have studied rounds versus communication tradeoffs for GT_n . Miltersen et al. show an $\Omega(n^{1/t} 2^{-O(t)})$ lower bound for t -round bounded error public coin randomised protocols for GT_n . Using our stronger round elimination lemma, we improve Miltersen et al.’s result.

Result 3. The bounded error public coin randomised t -round communication complexity of GT_n is lower bounded by $\Omega(n^{1/t} t^{-2})$. For bounded error quantum protocols with input-independent prior entanglement for GT_n , we have a lower bound of $\Omega(n^{1/t} t^{-1})$.

Remark. The lower bound for quantum protocols is better because, by definition, a quantum protocol always sends messages whose length is independent of the input.

There exists a bounded error classical randomised protocol for GT_n using t rounds of communication and having a complexity of $O(n^{1/t} \log n)$. Hence, for a constant number of rounds, our lower bound matches the upper bound to within logarithmic factors. For one round quantum protocols, our result implies an $\Omega(n)$ lower bound for GT_n (which is optimal to within constant factors), improving upon the previous $\Omega(n/\log n)$ lower bound of Klauck [21]. No rounds versus communication tradeoff for this problem, for more than one round, was known earlier in the quantum setting. If the number of rounds is unbounded, then there is a private coin classical randomised protocol for GT_n using $O(\log n)$ rounds of communication and having a complexity of $O(\log n)$ [32]. An $\Omega(\log n)$ lower bound for

the bounded error quantum communication complexity of GT_n (irrespective of the number of rounds) follows from Kremer's result [23] that the bounded error quantum communication complexity of a function is lower bounded (up to constant factors) by the logarithm of the one round (classical) deterministic communication complexity.

1.4. Our techniques

The starting point of our work is the paper of Miltersen, Nisan, Safra and Wigderson [29] showing lower bounds for randomised cell probe schemes for predecessor. The crux of Miltersen et al.'s lower bound is the following round elimination lemma for communication complexity.

Fact (Round elimination lemma, [29]). *Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ be a function. Let $\epsilon, \delta > 0$ be real numbers. Suppose that $\delta \leq \epsilon^2(100 \ln(8/\epsilon))^{-1}$. Suppose the communication game $f^{(n),A}$ has a $[t; l_1, \dots, l_t]^A$ public coin randomised protocol with error less than δ . Also suppose that $n \geq 20(l_1 \ln 2 + \ln 5)\epsilon^{-1}$. Then there is a $[t-1; l_2, \dots, l_t]^B$ public coin randomised protocol for f with error less than ϵ .*

Note that in the round elimination lemma of [29], the dependence between δ and ϵ is quadratic. In their paper, Miltersen et al. ask if their round elimination based approach can be strengthened to obtain Beame and Fich's [4] lower bound.

Our first observation is that if we can prove a stronger round elimination lemma in which δ and ϵ are related by a *small additive term*, and the additive term is upper bounded by $(\frac{1}{n})^{\Omega(1)}$, then we can obtain the lower bound of Beame and Fich for randomised cell probe schemes solving the predecessor problem.

Our next observation is that Klauck, Nayak, Ta-Shma and Zuckerman [22] have studied rounds versus quantum communication tradeoffs for the 'tree pointer chasing' problem using tools from quantum information theory. In fact, their quantum lower bound for the 'tree pointer chasing' problem is better than its previously known classical lower bound [29]! An important ingredient of their quantum lower bound was a quantum information-theoretic result called *average encoding* theorem. This result says informally that when messages carry very little information about the input, the average message is essentially as good as the individual messages. This result gives us a new way of attacking the round elimination problem. The information-theoretic round reduction arguments in [22] are average-case (under the uniform distribution on the inputs) arguments, and do not immediately give a worst-case result like the round elimination lemma. The information-theoretic arguments have to be combined with Yao's minimax lemma [43] (also used in the proof of the round elimination lemma in [29]) to prove the strong round elimination lemma of this paper. The information-theoretic approach brings out more clearly the intuition behind round elimination, as opposed to the ad hoc combinatorial proof in [29]. We believe that this strong round elimination lemma is an important technical contribution of this paper.

1.5. Subsequent work

After the paper [35] was published, Chakrabarti and Regev [6] proved an optimal lower bound for randomised cell probe schemes for the *approximate nearest neighbour* searching problem on the Hamming cube $\{0, 1\}^d$. Their lower bound is optimal in the same sense as ours, namely, if one wants a lower bound purely as a function of the universe size or the subset size, but not both. Their proof uses the classical version of our round elimination lemma to eliminate the messages of Bob who holds the stored subset. However, the round elimination lemma fails to eliminate the messages of Alice who holds the query point. In order to eliminate Alice's messages, Chakrabarti and Regev first compress her message by drawing on the message compression ideas of Jain, Radhakrishnan and Sen [17], and then remove it completely by using a simple but crucial *message switching* argument. Their message compression and switching argument better exploits the extreme asymmetry in the message lengths of Alice and Bob in this problem than our round elimination lemma alone. Though Chakrabarti and Regev proved their lower bound for classical cell probe schemes only, it is easy to see that their proof can be extended to address-only quantum cell probe schemes by using the later quantum message compression techniques of Jain, Radhakrishnan and Sen [18].

The main deficiency of our lower bounds for cell probe schemes for the predecessor problem is the following: our results only hold if we want lower bounds purely as a function of the universe size or the subset size, but not both together. In fact, the most general problem would be to give query bounds for predecessor optimal to within a

multiplicative universal constant that are simultaneously a function of the universe size, subset size, number of cells and the word length. Very recently, Pătraşcu and Thorup [33,34] solved this problem completely in a remarkable set of papers. For the upper bound, they refined the data structures of Beame and Fich [4], van Emde Boas et al. [39], fusion trees [12] and B-trees in order to get deterministic cell probe schemes tailored to different regimes of the parameters of the problem. Their lower bound holds for randomised query schemes also, and shows that their deterministic upper bounds are optimal in every regime of the parameters. To prove their lower bound when the word size is large, they combined our round elimination lemma with message compression and switching à la Chakrabarti and Regev [6]. To show their lower bound for small word size, they proved a novel *cell probe elimination lemma* using a direct sum based approach. This is their most important technical contribution, as communication complexity is powerless to prove the correct lower bounds when the word size is small. Their *cell probe elimination lemma* goes beyond communication complexity, and should find applications to other data structure problems in the future.

1.6. Organisation of the paper

We start with some preliminaries in the next section. Assuming the average encoding theorem, we prove an intermediate result in Section 3 which allows us to reduce the number of rounds of a communication protocol if the first message does not convey much information about the sender's input. Proofs of the classical and quantum versions of the average encoding theorem can be found in Appendix A for completeness. Using the intermediate result, we prove our strong round elimination lemma in Section 4. Sections 3 and 4 each have two subsections: the first one treats the classical version of the results and the second one treats the quantum version. Using the strong round elimination lemma, we prove the optimal lower bound for the predecessor problem in Section 5. The rounds versus communication tradeoff for the 'greater-than' problem is sketched in Section 6. We end with a discussion of our work and indicate some open problems in Section 7.

2. Preliminaries

2.1. The address-only quantum cell probe model

In this subsection, we recall the definition of the *address-only quantum cell probe model* from [36]. A quantum (s, w, t) *cell probe scheme* for a static data structure problem $f: D \times Q \rightarrow A$ has two components: a classical deterministic *storage scheme* that stores the data $d \in D$ in a table $T[d]$ using s cells each containing w bits, and a quantum *query scheme* that answers queries by 'quantumly probing a cell at a time' t times. Formally speaking, the table $T[d]$ is made available to the query algorithm in the form of an oracle unitary transform O_d . To define O_d formally, we represent the basis states of the query algorithm as $|j, b, z\rangle$, where $j \in [s-1]$ is a binary string of length $\log s$, b is a binary string of length w , and z is a binary string of some fixed length. Here, j denotes the address of a cell in the table $T[d]$, b denotes the qubits which will hold the contents of a cell and z stands for the rest of the qubits ('work qubits') in the query algorithm. O_d maps $|j, b, z\rangle$ to $|j, b \oplus T[d]_j, z\rangle$, where $T[d]_j$ is a bit string of length w and denotes the contents of the j th cell in $T[d]$. A quantum query scheme with t probes is just a sequence of unitary transformations

$$U_0 \rightarrow O_d \rightarrow U_1 \rightarrow O_d \rightarrow \cdots \rightarrow U_{t-1} \rightarrow O_d \rightarrow U_t$$

where U_j 's are arbitrary unitary transformations that do not depend on d (U_j 's represent the internal computations of the query algorithm). For a query $q \in Q$, the computation starts in a computational basis state $|q\rangle|0\rangle$, where we assume that the ancilla qubits are initially in the basis state $|0\rangle$. Then we apply in succession, the operators $U_0, O_d, U_1, \dots, U_{t-1}, O_d, U_t$, and measure the final state. The answer consists of the values on some of the output wires of the circuit. We say that the scheme has worst case error probability less than ϵ if the answer is equal to $f(d, q)$, for every $(d, q) \in D \times Q$, with probability greater than $1 - \epsilon$. The term 'bounded error quantum scheme' means that $\epsilon = 1/3$.

We can now formally define the *address-only quantum cell probe model*. Here the storage scheme is classical deterministic as before, but the query scheme is restricted to be 'address-only quantum.' This means that the state vector before a query to the oracle O_d is always a *tensor product* of a state vector on the address and work qubits (the (j, z) part in (j, b, z) above), and a state vector on the data qubits (the b part in (j, b, z) above). The state vector on the data qubits before a query to the oracle O_d is *independent of the query element q and the data d* but can vary

with the probe number. Intuitively, we are only making use of quantum parallelism over the address lines of a query. This mode of querying a table subsumes classical (deterministic or randomised) querying, and also many non-trivial quantum algorithms like Grover’s algorithm [14], Farhi et al.’s algorithm [10], Høyer et al.’s algorithm [15] etc. satisfy the ‘address-only’ condition. For classical (deterministic or randomised) querying, the state vector on the data qubits is $|0\rangle$, independent of the probe number. For Grover’s algorithm and Farhi et al.’s algorithm, the state vector on the data qubit is $(|0\rangle - |1\rangle)/\sqrt{2}$, independent of the probe number. For Høyer et al.’s algorithm, the state vector on the data qubit is $|0\rangle$ for some probe numbers, and $(|0\rangle - |1\rangle)/\sqrt{2}$ for the other probe numbers.

2.2. Quantum communication protocols

In this paper, we adopt the ‘interacting unitary quantum circuits’ definition of quantum communication protocols of Yao [42]. Thus, Alice and Bob send a certain number of fixed length messages to each other, and the number and length of these messages is independent of their inputs. If Alice’s and Bob’s inputs are in computational basis states, the global state of all the qubits of Alice and Bob is pure at all times during the execution of the protocol. Measurements are not allowed during the execution of the protocol. At the end of the protocol, the last recipient of a message makes a von Neumann measurement in the computational basis of certain qubits (the ‘answer qubits’) in her possession in order to determine the answer of the protocol. The choice of ‘answer qubits’ is independent of Alice’s and Bob’s inputs.

We require that Alice and Bob make a secure copy of their inputs before beginning the quantum communication protocol. This is possible since the inputs to Alice and Bob are in computational basis states e.g. CNOT gates can be used for this purpose. Thus, without loss of generality, the input qubits of Alice and Bob are never sent as messages, their state remains unchanged throughout the protocol, and they are never measured, i.e. some work qubits are measured to determine the result of the protocol. We call such quantum protocols *secure* and will assume henceforth that all our quantum protocols are secure.

We now define the concept of a *safe* quantum communication protocol, which will be used in the statement of the quantum round elimination lemma.

Definition 4 (Safe quantum protocol). A $[t; c; l_1, \dots, l_t]^A$ ($[t; c; l_1, \dots, l_t]^B$) safe quantum communication protocol is a secure quantum protocol where Alice (Bob) starts the communication, the first message is $l_1 + c$ qubits long, the i th message, for $i \geq 2$, is l_i qubits long, and the communication goes on for t rounds. We think of the first message as having two parts: the ‘main part’ which is l_1 qubits long, and the ‘safe overhead part’ which is c qubits long. The density matrix of the ‘safe overhead’ is independent of the inputs to Alice and Bob.

Remarks. (1) The safe overhead is nothing but a way to send input independent prior entanglement from Alice to Bob. The reason we use this notation is that we will later define safe public coin quantum protocols where there will be two kinds of input independent prior entanglement, the first being the one provided by the safe overhead and the second being the one provided by the public coin.

(2) The reason for defining the concept of a safe overhead, intuitively speaking, is as follows. The communication games arising from data structure problems often have an asymmetry between the message lengths of Alice and Bob. This asymmetry is crucial to prove lower bounds on the number of rounds of communication. In the previous quantum round reduction arguments (e.g. those of Klauck et al. [22]), the complexity of the first message in the protocol increases quickly as the number of rounds is reduced and the asymmetry gets lost. This leads to a problem where the first message soon gets big enough to potentially convey substantial information about the input of one player to the other, destroying any hope of proving strong lower bounds on the number of rounds. The concept of a safe protocol allows us to get around this problem. We show through a careful quantum information theoretic analysis of the round reduction process, that in a safe protocol, though the complexity of the first message increases a lot, this increase is confined to the safe overhead and so, the information content does not increase much. This is the key property which allows us to prove a round elimination lemma for safe quantum protocols.

In this paper we will deal with quantum protocols with *public coins*. Intuitively, a public coin quantum protocol is a probability distribution over finitely many (*coinless*) quantum protocols. We shall henceforth call the standard definition of a quantum protocol without prior entanglement as *coinless*. Our definition is similar to the classical

scenario, where a randomised protocol with public coins is a probability distribution over finitely many deterministic protocols. We note however, that our definition of a public coin quantum protocol is *not* the same as that of a quantum protocol with prior entanglement, which has been studied previously (see e.g. [7]). Our definition is weaker, in that it does not allow the unitary transformations of Alice and Bob to alter the ‘public coin.’

Definition 5 (*Public coin quantum protocol*). In a quantum protocol with a public coin, there is, before the start of the protocol, a quantum state called a *public coin*, of the form $\sum_c \sqrt{p_c} |c\rangle_A |c\rangle_B$, where the subscripts denote ownership of qubits by Alice and Bob, p_c are finitely many non-negative real numbers and $\sum_c p_c = 1$. Alice and Bob make (entangled) copies of their respective halves of the public coin using CNOT gates before commencing the protocol. The unitary transformations of Alice and Bob during the protocol do not touch the public coin. The public coin is never measured, nor is it ever sent as a message.

Hence, one can think of the public coin quantum protocol to be a probability distribution, with probability p_c , over finitely many coinless quantum protocols indexed by the coin basis states $|c\rangle$. A *safe public coin* quantum protocol is similarly defined as a probability distribution over finitely many safe coinless quantum protocols.

Remarks. (1) We need to define public coin quantum protocols in order to make use of the harder direction of Yao’s minimax lemma [43]. The minimax lemma is the main tool which allows us to convert ‘average case’ round reduction arguments to ‘worst case’ arguments. We need ‘worst case’ round reduction arguments in proving lower bounds for the rounds complexity of communication games arising from data structure problems. This is because many of these lower bound proofs use some notion of “self-reducibility,” arising from the original data structure problem, which fails to hold in the ‘average case’ but holds for the ‘worst case.’ The quantum round reduction arguments of Klauck et al. [22] are ‘average case’ arguments, and this is one of the reasons why they do not suffice to prove lower bounds for the rounds complexity of communication games arising from data structure problems.

(2) For $[t; c; l_1, \dots, l_t]^A$ safe quantum protocols computing a function f , Yao’s minimax lemma says that the infimum (worst-case) error of a public coin protocol is equal to the supremum over all input probability distributions of the infimum distributional error of a coinless protocol.

2.3. Predecessor searching and communication complexity

We first describe the connection between the address-only quantum cell probe complexity of a static data structure problem and the quantum communication complexity of an associated communication game. Let $f : D \times Q \rightarrow A$ be a static data structure problem. Consider a two-party communication problem where Alice is given a query $q \in Q$, Bob is given a datum $d \in D$, and they have to communicate and find out the answer $f(d, q)$. We have the following lemma, which is a quantum analogue of a lemma of Miltersen [28] relating cell probe complexity to communication complexity in the classical setting.

Lemma 1. *Suppose we have an (s, w, t) quantum cell probe solution to a static data structure problem $f : D \times Q \rightarrow A$. Then we have a $(2t, 0, \log s + w, \log s + w)^A$ safe coinless quantum protocol for the corresponding communication problem. If the query scheme is address-only, we also have a $(2t, 0, \log s, \log s + w)^A$ safe coinless quantum protocol for the corresponding communication problem. The error probability of the communication protocol is the same as that of the cell probe scheme.*

Proof. Given a quantum (s, w, t) cell probe solution to the static data structure problem f , we can get a $(2t, 0, \log s + w, \log s + w)^A$ safe coinless quantum protocol for the corresponding communication problem by just simulating the cell probe solution. If in addition, the query scheme is address-only, the messages from Alice to Bob need consist only of the ‘address’ part. This can be seen as follows. Let the state vector of the data qubits before the i th query be $|\theta_i\rangle$. $|\theta_i\rangle$ is independent of the query element and the stored data. Bob keeps t special ancilla registers in states $|\theta_i\rangle$, $1 \leq i \leq t$, at the start of the protocol P . These special ancilla registers are in tensor with the rest of the qubits of Alice and Bob at the start of P . Protocol P simulates the cell probe solution, but with the following modification. To simulate the i th query of the cell probe solution, Alice prepares her ‘address’ and ‘data’ qubits as in the query scheme, but sends the ‘address’ qubits only. Bob treats those ‘address’ qubits together with $|\theta_i\rangle$ in the i th special ancilla register

as Alice’s query, and performs the oracle table transformation on them. He then sends these qubits (both the ‘address’ as well as the i th special register qubits) to Alice. Alice exchanges the contents of the i th special register with her ‘data’ qubits (i.e. exchanges the basis states), and proceeds with the simulation of the query scheme. This gives us a $(2t, 0, \log s, \log s + w)^A$ safe coinless quantum protocol with the same error probability as that of the cell probe query scheme. \square

Remark. In many natural data structure problems $\log s$ is much smaller than w and thus, in the address-only quantum case, we get a $(2t, 0, \log s, O(w))^A$ safe protocol. In the classical setting of [28], one gets a $(2t, 0, \log s, w)^A$ protocol. This asymmetry in message lengths is crucial in proving non-trivial lower bounds on t . The concept of a safe quantum protocol helps us in exploiting this asymmetry.

We now recall some facts about the connection between cell probe schemes for predecessor and communication complexity of rank parity from [29]. We give proof sketches of these facts for completeness.

Definition 6. A $(t, a, b)^A ((t, a, b)^B)$ classical communication protocol is a $[t; l_1, \dots, l_t]^A ([t; l_1, \dots, l_t]^B)$ classical protocol, where assuming Alice (Bob) starts, $l_i = a$ for i odd and $l_i = b$ for i even ($l_i = b$ for i odd and $l_i = a$ for i even). A $(t, c, a, b)^A ((t, c, a, b)^B)$ quantum communication protocol is a $[t; c; l_1, \dots, l_t]^A ([t; c; l_1, \dots, l_t]^B)$ safe quantum protocol, where assuming Alice (Bob) starts, $l_i = a$ for i odd and $l_i = b$ for i even ($l_i = b$ for i odd and $l_i = a$ for i even).

Definition 7 (Rank parity). In the rank parity communication game $\text{PAR}_{p,q}$, Alice is given a bit string x of length p , Bob is given a set S of bit strings of length p , $|S| \leq q$, and they have to communicate and decide whether the rank of x in S (treating the bit strings as integers) is odd or even. By the rank of x in S , we mean the cardinality of the set $\{y \in S \mid y \leq x\}$. In the communication game $\text{PAR}_{p,q}^{(k),A}$, Alice is given k bit strings x_1, \dots, x_k each of length p , Bob is given a set S of bit strings of length p , $|S| \leq q$, an index $i \in [k]$, and copies of x_1, \dots, x_{i-1} ; they have to communicate and decide whether the rank of x_i in S is odd or even. In the communication game $\text{PAR}_{p,q}^{(k),B}$, Alice is given a bit string x of length p and an index $i \in [k]$, Bob is given k sets S_1, \dots, S_k of bit strings of length p , $|S_j| \leq q$, $1 \leq j \leq k$; they have to communicate and decide whether the rank of x in S_i is odd or even.

Fact 1. Let m be a positive integer such that m is a power of 2. Suppose that there is a $(n^{O(1)}, (\log m)^{O(1)}, t)$ randomised (address-only quantum) cell probe scheme for the (m, n) -static predecessor problem. Then the rank parity communication game $\text{PAR}_{\log m, n}$ has a $(2t + O(1), O(\log n), (\log m)^{O(1)})^A$ private coin randomised (safe coinless quantum) protocol. The error probability of the communication protocol is the same as that of the cell probe scheme.

Proof. Consider the static rank parity data structure problem where the storage scheme has to store a set $S \subseteq [m]$, $|S| \leq n$, and the query scheme, given a query $x \in [m]$, has to decide whether the rank of x in S is odd or even. Fredman, Komlós and Szemerédi [11] have shown the existence of two-level perfect hash tables containing, for each member y of the stored subset S , y ’s rank in S , and using $O(n)$ cells of word size $O(\log m)$ and requiring only $O(1)$ deterministic cell probes. Combining a $(n^{O(1)}, (\log m)^{O(1)}, t)$ cell probe solution to the static predecessor problem with such a perfect hash table gives us a $(n^{O(1)} + O(n), \max((\log m)^{O(1)}, O(\log m)), t + O(1))$ cell probe solution to the static rank parity problem. The error probability of the cell probe scheme for the rank parity problem is the same as the error probability of the cell probe scheme for the predecessor problem. Converting the cell probe scheme for the rank parity problem into a communication protocol by [28] (by Lemma 1), we get a $(2t + O(1), O(\log n), (\log m)^{O(1)})^A$ private coin randomised (safe coinless quantum) protocol for the rank parity communication game $\text{PAR}_{\log m, n}$. The error probability of the communication protocol is the same as that of the cell probe scheme for the predecessor problem. \square

Fact 2. Let k, p be positive integers such that $k \mid p$. A communication protocol with Alice starting for $\text{PAR}_{p,q}$ gives us a communication protocol with Alice starting for $\text{PAR}_{p/k,q}^{(k),A}$ with the same message complexity, number of rounds and error probability.

Proof. Consider the problem $\text{PAR}_{p/k,q}^{(k),A}$. Alice, who is given x_1, \dots, x_k , computes the concatenation $\hat{x} \triangleq x_1 \cdot x_2 \cdots x_k$. Bob, who is given S, i and x_1, \dots, x_{i-1} , computes

$$\hat{S} \triangleq \{x_1 \cdot x_2 \cdots x_{i-1} \cdot y \cdot 0^{p(1-i/k)} : y \in S\}.$$

After this, Alice and Bob run the protocol for $\text{PAR}_{p,q}$ on inputs \hat{x}, \hat{S} to solve the problem $\text{PAR}_{p/k,q}^{(k),A}$. \square

Fact 3. Let k, q be positive integers such that $k \mid q$ and k is a power of 2. A communication protocol with Bob starting for $\text{PAR}_{p,q}$ gives us a communication protocol with Bob starting for $\text{PAR}_{p-\log k-1,q/k}^{(k),B}$ with the same message complexity, number of rounds and error probability.

Proof. Consider the problem $\text{PAR}_{p-\log k-1,q/k}^{(k),B}$. Alice, who is given x and i , computes the concatenation $\hat{x} \triangleq (i-1) \cdot 0 \cdot x$. Bob, who is given S_1, \dots, S_k , computes the sets S'_1, \dots, S'_k where

$$S'_j \triangleq \begin{cases} \{(j-1) \cdot 0 \cdot y : y \in S_j\} & \text{if } |S_j| \text{ is even,} \\ \{(j-1) \cdot 0 \cdot y : y \in S_j\} \cup \{(j-1) \cdot 1^{p-\log k}\} & \text{if } |S_j| \text{ is odd.} \end{cases}$$

Above, the integers $(i-1), (j-1)$ are to be thought of as bit strings of length $\log k$. Bob also computes $\hat{S} \triangleq \bigcup_{j=1}^k S'_j$. Alice and Bob then run the protocol for $\text{PAR}_{p,q}$ on inputs \hat{x}, \hat{S} to solve the problem $\text{PAR}_{p-\log k-1,q/k}^{(k),B}$. \square

2.4. Some classical information theoretic facts

In this subsection, we discuss some classical information theoretic facts which will be used in the proof of our improved classical round elimination lemma. For a good account of classical information theory, see the book by Cover and Thomas [8].

In this paper, all random variables have finite range and all sample spaces have finite cardinality. Let X, Y, Z be random variables with some joint distribution. The *Shannon entropy* of X is defined as $H(X) \triangleq -\sum_x \Pr[X=x] \log \Pr[X=x]$. The *mutual information* of X and Y is defined as $I(X:Y) \triangleq H(X) + H(Y) - H(XY)$. If the range of X has cardinality at most d , $I(X:Y) \leq \log d$. $I(X:Y) = 0$ iff X and Y are independent. Let $I((X:Y) | Z=z)$ denote the mutual information of X and Y conditioned on the event $Z=z$.

The next fact follows easily from the definitions.

Fact 4. Let X, Y, Z be random variables with some joint distribution. Then,

- (a) $I(XY:Z) = I(X:Z) + I(Y:ZX) - I(X:Y)$. In particular, if X, Y are independent $I(XY:Z) = I(X:Z) + I(Y:ZX)$.
- (b) $I(Y:ZX) = I(X:Y) + \mathbf{E}_X[I((Y:Z) | X=x)]$, where is expectation is over the marginal distribution of X .

We use total variation distance to quantify the distance between two probability distributions.

Definition 8 (Total variation distance). Let P, Q be probability distributions on the same sample space Ω . The *total variation distance* (also known as the ℓ_1 -distance) between P and Q , denoted by $\|P - Q\|_1$, is defined as $\|P - Q\|_1 \triangleq \sum_{x \in \Omega} |P(x) - Q(x)|$.

We will need the *average encoding* theorem of Klauck, Nayak, Ta-Shma and Zuckerman [22]. Klauck et al. actually prove a quantum version of this theorem in their paper, but we will use the classical version in the proof of our round elimination lemma. Intuitively speaking, the theorem says that if the mutual information between a random variable and its randomised encoding is small, then the probability distributions on the code words for various values of the random variable are close to the average probability distribution on the code words.

Fact 5 (Average encoding theorem, [22]). Let X, M be correlated random variables. Let p_x denote the (marginal) probability that $X = x$, and Π^x denote the conditional distribution of M given that $X = x$. Let Π denote the (marginal) probability distribution of M , i.e. $\Pi = \sum_x p_x \Pi^x$. Then,

$$\sum_x p_x \|\Pi^x - \Pi\|_1 \leq \sqrt{(2 \ln 2) I(X : M)}.$$

A self-contained classical proof, without using quantum information theory, of this fact can be found in Appendix A for completeness.

2.5. Some quantum information theoretic facts

In this subsection, we discuss some quantum information theoretic facts which will be used in the proof of our quantum round elimination lemma. For a good account of quantum information theory, see the book by Nielsen and Chuang [31].

In this paper, all quantum systems, Hilbert spaces and superoperators are finite-dimensional. Let X, Y be quantum systems with some joint density matrix ρ_{XY} . If ρ_X is the reduced density matrix of X , the *von Neumann entropy* of X is defined as $S(X) \triangleq -\text{Tr} \rho_X \log \rho_X$. The *mutual information* of X and Y is defined as $I(X : Y) \triangleq S(X) + S(Y) - S(XY)$. If X is at most d -dimensional, $I(X : Y) \leq 2 \log d$. $I(X : Y) = 0$ iff X and Y are independent, i.e. $\rho_{XY} = \rho_X \otimes \rho_Y$. Suppose X, Y, Z are quantum systems with some joint density matrix where Z is a classical random variable, i.e. the reduced density matrix of Z is diagonal in the computational basis. Let $I((X : Y) | Z = z)$ denote the mutual information of X and Y conditioned on the event $Z = z$.

The next fact follows easily from the definitions.

Fact 6. Let X, Y, Z, W be quantum systems with some joint density matrix, where X and Y are classical random variables. Then,

- (a) $I(XY : Z) = I(X : Z) + I(Y : ZX) - I(X : Y)$. In particular, if X, Y are independent $I(XY : Z) = I(X : Z) + I(Y : ZX)$.
- (b) $I(Y : ZX) = I(X : Y) + \mathbf{E}_X[I((Y : Z) | X = x)]$, where \mathbf{E}_X is expectation is over the marginal distribution of X .
- (c) Suppose W is independent of X and Z is supported on m qubits. Then, $I(X : ZW) \leq 2m$.

Remarks. (1) Fact 6(c) is the key observation allowing us to “ignore” the size of the “safe” overhead W in quantum round elimination applications. In these applications, the complexity of the first message in the protocol increases quickly, but the blow up is confined to the “safe” overhead. Earlier round reduction arguments were unable to handle this large blow up in the complexity of the first message.

(2) In Fact 6(c), if Z is a classical random variable, we get the improved inequality $I(X : ZW) \leq m$.

We use trace distance to quantify the distance between two density matrices.

Definition 9 (Trace distance). Let ρ, σ be density matrices in the same Hilbert space. The *trace distance* between ρ and σ , denoted by $\|\rho - \sigma\|_1$, is defined as $\|\rho - \sigma\|_1 \triangleq \text{Tr} \sqrt{(\rho - \sigma)^\dagger (\rho - \sigma)}$.

If ρ is a density matrix in a Hilbert space \mathcal{H} and \mathcal{M} is a general measurement, i.e. a POVM on \mathcal{H} , let $\mathcal{M}\rho$ denote the probability distribution on the (classical) outcomes of \mathcal{M} got by performing measurement \mathcal{M} on ρ . The importance of the trace distance as a metric on density matrices stems from the following fundamental fact (see e.g. [1]).

Fact 7. Let ρ_1, ρ_2 be two density matrices in the same Hilbert space \mathcal{H} . Let \mathcal{M} be a POVM on \mathcal{H} . Then, $\|\mathcal{M}\rho_1 - \mathcal{M}\rho_2\|_1 \leq \|\rho_1 - \rho_2\|_1$.

Let \mathcal{H}, \mathcal{K} be disjoint Hilbert spaces such that $\dim(\mathcal{K}) \geq \dim(\mathcal{H})$. Let ρ be a density matrix in \mathcal{H} and $|\psi\rangle$ be a pure state in $\mathcal{H} \otimes \mathcal{K}$. $|\psi\rangle$ is said to be a *purification* of ρ if $\text{Tr}_{\mathcal{K}} |\psi\rangle\langle\psi| = \rho$. We will require the following basic fact about two purifications of the same density matrix.

Fact 8. Let \mathcal{H}, \mathcal{K} be disjoint Hilbert spaces such that $\dim(\mathcal{K}) \geq \dim(\mathcal{H})$. Let ρ be a density matrix in \mathcal{H} and $|\psi\rangle, |\phi\rangle$ be two purifications of ρ in $\mathcal{H} \otimes \mathcal{K}$. Then there is a local unitary transformation U on \mathcal{K} such that $|\psi\rangle = (I \otimes U)|\phi\rangle$, where I is the identity operator on \mathcal{H} .

We now state an improved version of the quantum *average encoding* theorem of Klauck, Nayak, Ta-Shma and Zuckerman [22]. This improved version follows from a direct connection between fidelity and relative entropy described in [9], and was also independently observed by Klauck (private communication). Intuitively speaking, the theorem says that if the mutual information between a random variable and its quantum encoding is small, then given any purifications of the quantum code words for various values of the random variable, one can find purifications of the average code word that are close to the respective purifications of the quantum code words.

Fact 9. Let X, M be quantum systems with some joint density matrix. Let X be a classical random variable and p_x denote the (marginal) probability that $X = x$. Let \mathcal{H} denote the Hilbert space of M and ρ^x denote the conditional density matrix of M given that $X = x$. Let ρ denote the reduced density matrix of M . Note that $\rho = \sum_x p_x \rho^x$. Let \mathcal{K} denote a disjoint Hilbert space such that $\dim(\mathcal{K}) \geq \dim(\mathcal{H})$. Let $|\psi^x\rangle$ be purifications of ρ^x in $\mathcal{H} \otimes \mathcal{K}$. Then there exist purifications $|\phi^x\rangle$ of ρ in $\mathcal{H} \otimes \mathcal{K}$ such that

$$\sum_x p_x \left\| |\psi^x\rangle\langle\psi^x| - |\phi^x\rangle\langle\phi^x| \right\|_1 \leq \sqrt{(4 \ln 2) I(X : M)}.$$

A proof of this fact can be found in Appendix A for completeness.

3. Reducing the number of rounds

In this section, we prove an intermediate result which will be required to prove our strong round elimination lemmas. In Section 3.1 we prove the intermediate result in its classical version (Lemma 2), whereas in Section 3.2 we prove the intermediate result in its quantum version (Lemma 3). The proof of Lemma 2 is similar to the proof of Lemma 4.4 in [22] (see also [27]), but much simpler since we are in the classical setting. The proof of Lemma 3 is a slight refinement of the proof of Lemma 4.4 in [22], and is included here for completeness. Intuitively speaking, the intermediate result says that if the first message in a communication protocol carries little information about the sender’s input, under some probability distribution on Alice’s and Bob’s inputs, then it can be eliminated, giving rise to a protocol where the other player starts, with one less round of communication, smaller message complexity, and with similar average error probability with respect to the same probability distribution on Alice’s and Bob’s inputs.

Consider a communication protocol \mathcal{P} computing a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$. For an input $(x, y) \in \mathcal{X} \times \mathcal{Y}$, we define the error $\epsilon_{x,y}^{\mathcal{P}}$ of \mathcal{P} on (x, y) to be the probability that the result of \mathcal{P} on input (x, y) is not equal to $f(x, y)$. For a protocol \mathcal{P} , given a probability distribution D on $\mathcal{X} \times \mathcal{Y}$, we define the average error $\epsilon_D^{\mathcal{P}}$ of \mathcal{P} with respect to D as the expectation over D of the error of \mathcal{P} on inputs $(x, y) \in \mathcal{X} \times \mathcal{Y}$. We define $\epsilon^{\mathcal{P}}$ to be the maximum error of \mathcal{P} on inputs $(x, y) \in \mathcal{X} \times \mathcal{Y}$, i.e. $\epsilon^{\mathcal{P}}$ is the error of protocol \mathcal{P} .

3.1. Classical

Lemma 2. Suppose $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ is a function. Let D be a probability distribution on $\mathcal{X} \times \mathcal{Y}$, and \mathcal{P} be a $[t; l_1, \dots, l_t]^A$ private coin randomised protocol for f . Let X stand for the classical random variable denoting Alice’s input, M stand for the random variable denoting the first message of Alice in \mathcal{P} , and $I(X : M)$ denote the mutual information between X and M when the inputs to \mathcal{P} are distributed according to D . Then there exists a $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol \mathcal{Q} for f , such that $\epsilon_D^{\mathcal{Q}} \leq \epsilon_D^{\mathcal{P}} + (1/2)((2 \ln 2)I(X : M))^{1/2}$.

Proof. We first give an overview of the idea of the proof before getting down to the details. The proof proceeds in stages.

Idea of Stage 1. Starting from protocol \mathcal{P} , we construct a $[t; l_1, \dots, l_t]^A$ private coin protocol \mathcal{P}' where the first message is independent of Alice’s input, and $\epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((2 \ln 2)I(X : M))^{1/2}$. The important idea here is to first generate Alice’s message using a new private coin without ‘looking’ at her input, and after that, to adjust Alice’s old private coin in a suitable manner so as to be consistent with her message and input.

Idea of Stage 2. Since the first message of \mathcal{P}' is independent of Alice's input, Bob can generate it himself. Doing this and setting coin tosses appropriately gives us a $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol \mathcal{Q} for f such that $\epsilon_D^{\mathcal{Q}} \leq \epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((2 \ln 2)I(X : M))^{1/2}$.

We now give the details of the proof. Let Π^x be the probability distribution of the first message M of protocol \mathcal{P} when Alice's input $X = x$. Define $\Pi \triangleq \sum_x d_x \Pi^x$, where d_x is the marginal probability of $X = x$ under distribution D . Π is the probability distribution of the average first message of \mathcal{P} under distribution D . For $x \in \mathcal{X}$ and an instance m of the first message of Alice, let q_r^{xm} denote the conditional probability that the private coin toss of Alice results in r , given that Alice's input is x and her first message in \mathcal{P} is m . If in \mathcal{P} message m cannot occur when Alice's input is x , then we define $q_r^{xm} \triangleq 0$. Let π_m^x denote the probability that the first message of Alice in \mathcal{P} is m , given that her input is x . Let π_m denote the probability that the first message of Alice in \mathcal{P} is m , when Alice's and Bob's inputs are distributed according to D . Then, $\pi_m = \sum_x d_x \pi_m^x$.

Stage 1. We construct a $[t; l_1, \dots, l_t]^A$ private coin randomised protocol \mathcal{P}' for f with average error under distribution D $\epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((2 \ln 2)I(X : M))^{1/2}$, and where the probability distribution of the first message is independent of the input to Alice. We now describe the protocol \mathcal{P}' . Suppose Alice is given $x \in \mathcal{X}$ and Bob is given $y \in \mathcal{Y}$. Alice tosses a fresh private coin to pick m with probability π_m and set her old private coin to r with probability q_r^{xm} . After this, Alice and Bob behave as in protocol \mathcal{P} (henceforth, Alice ignores the new private coin which she had tossed to generate her first message m). Hence in \mathcal{P}' the probability distribution of the first message is independent of Alice's input.

Let us now compare the situations in protocols \mathcal{P} and \mathcal{P}' when Alice's input is x , Bob's input is y , Alice has finished tossing her private coins (both old and new), but no communication has taken place as yet. In protocol \mathcal{P} , the probability that Alice's private coin toss results in r is $\sum_m \pi_m^x q_r^{xm}$. In protocol \mathcal{P}' , the probability that Alice's old private coin is set to r is $\sum_m \pi_m q_r^{xm}$. Thus, the total variation distance between the probability distributions on Alice's old private coin is

$$\begin{aligned} \sum_r \left| \sum_m q_r^{xm} (\pi_m^x - \pi_m) \right| &\leq \sum_r \sum_m q_r^{xm} |\pi_m^x - \pi_m| = \sum_m \left(|\pi_m^x - \pi_m| \sum_r q_r^{xm} \right) = \sum_m |\pi_m^x - \pi_m| \\ &= \|\Pi^x - \Pi\|_1. \end{aligned}$$

Hence, the error probability of \mathcal{P}' on input x, y $\epsilon_{x,y}^{\mathcal{P}'} \leq \epsilon_{x,y}^{\mathcal{P}} + (1/2)\|\Pi^x - \Pi\|_1$. Let d_{xy} be the probability that $(X, Y) = (x, y)$ under distribution D . Then, the average error of \mathcal{P}' under distribution D

$$\begin{aligned} \epsilon_D^{\mathcal{P}'} &= \sum_{x,y} d_{xy} \epsilon_{x,y}^{\mathcal{P}'} \leq \sum_{x,y} d_{xy} \left(\epsilon_{x,y}^{\mathcal{P}} + \frac{1}{2} \|\Pi^x - \Pi\|_1 \right) = \epsilon_D^{\mathcal{P}} + \frac{1}{2} \sum_x d_x \|\Pi^x - \Pi\|_1 \\ &\leq \epsilon_D^{\mathcal{P}} + \frac{1}{2} ((2 \ln 2)I(X : M))^{1/2}. \end{aligned}$$

The last inequality follows from Fact 5.

Stage 2. We now construct our desired $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol \mathcal{Q} for f with $\epsilon_D^{\mathcal{Q}} \leq \epsilon_D^{\mathcal{P}'}$. Suppose all the coin tosses of Alice and Bob in \mathcal{P}' are done publicly before any communication takes place. Now there is no need for the first message from Alice to Bob, because Bob can reconstruct the message by looking at the public coin tosses. This gives us a $[t - 1; l_2, \dots, l_t]^B$ public coin protocol \mathcal{Q}' such that $\epsilon_{x,y}^{\mathcal{Q}'} = \epsilon_{x,y}^{\mathcal{P}'}$ for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$. By setting the public coin tosses of \mathcal{Q}' to an appropriate value, we get a $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol \mathcal{Q} such that $\epsilon_D^{\mathcal{Q}} \leq \epsilon_D^{\mathcal{Q}'} = \epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((2 \ln 2)I(X : M))^{1/2}$.

This completes the proof of Lemma 2. \square

3.2. Quantum

Lemma 3. Suppose $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ is a function. Let D be a probability distribution on $\mathcal{X} \times \mathcal{Y}$, and \mathcal{P} be a $[t; c; l_1, \dots, l_t]^A$ safe coinless quantum protocol for f . Let X stand for the classical random variable denoting Alice's input, M denote the first message of Alice in \mathcal{P} , and $I(X : M)$ denote the mutual information between X and M when the inputs to \mathcal{P} are distributed according to D . Then there exists a $[t - 1; c + l_1; l_2, \dots, l_t]^B$ safe coinless quantum protocol \mathcal{Q} for f , such that $\epsilon_D^{\mathcal{Q}} \leq \epsilon_D^{\mathcal{P}} + (1/2)((4 \ln 2)I(X : M))^{1/2}$.

Proof. We first give an overview of the plan of the proof, before getting down to the details. The proof proceeds in stages. Stage 1 of the quantum proof corresponds to Stage 1 of the classical proof, and Stages 2A and 2B of the quantum proof together correspond to Stage 2 of the classical proof.

Idea of Stage 1. Starting from protocol \mathcal{P} , we construct a $[t; c; l_1, \dots, l_t]^A$ safe coinless quantum protocol \mathcal{P}' where the first message is independent of Alice’s input, and $\epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((4 \ln 2)I(X : M))^{1/2}$. The important idea here is to generate a purification $|\phi^x\rangle$ of the average first message ρ of Alice in \mathcal{P} that is close to the purification $|\psi^x\rangle$ of the message ρ^x of Alice in \mathcal{P} when her input is x . The existence of such a purification $|\phi^x\rangle$ is guaranteed by Fact 9.

Idea of Stage 2A. Since the first message of \mathcal{P}' is independent of Alice’s input (in fact its density matrix is ρ), Bob can generate it himself. This suffices if \mathcal{P}' is a one-round protocol, and completes the proof of Lemma 3 for such protocols. But if \mathcal{P}' has more than one round, it is also necessary for Bob to achieve the correct entanglement between Alice’s work qubits and the first message, i.e. it is necessary that the joint state of Alice’s work qubits and the qubits of the first message be $|\phi^x\rangle$. Bob achieves this by first sending a safe message of $l_1 + c$ qubits. If Alice’s input is x , she then applies a unitary transformation V_x on her work qubits in order to make the joint state of her work qubits and the qubits of the first message $|\phi^x\rangle$. The existence of such a V_x follows from Fact 8. Doing all this gives us a $[t + 1; c + l_1; 0, 0, l_2, \dots, l_t]^B$ safe coinless quantum protocol \mathcal{Q}' for f such that $\epsilon_{x,y}^{\mathcal{Q}'} = \epsilon_{x,y}^{\mathcal{P}'}$ for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$.

Idea of Stage 2B. Since the first message of Alice in \mathcal{Q}' is zero qubits long, Bob can concatenate his first two messages, giving us a $[t - 1; c + l_1; l_2, \dots, l_t]^B$ safe coinless quantum protocol \mathcal{Q} for f such that $\epsilon_{x,y}^{\mathcal{Q}} = \epsilon_{x,y}^{\mathcal{Q}'}$ for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$. The technical reason behind this is that unitary transformations on disjoint sets of qubits commute. Note that $\epsilon_D^{\mathcal{Q}} = \epsilon_D^{\mathcal{Q}'} = \epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((4 \ln 2)I(X : M))^{1/2}$.

We now give the details of the proof. Let ρ^x be the density matrix of the first message M of protocol \mathcal{P} when Alice’s input X is x . Let A and B denote Alice’s work qubits excluding the qubits of M and Bob’s work qubits respectively. Let $|\psi^x\rangle_{AM}$ be the joint pure state of AM in \mathcal{P} when Alice’s input X is x , she has finished preparing her first message, but no communication has taken place as yet. Without loss of generality, the number of qubits in A is at least the number of qubits in M . Define $\rho \triangleq \sum_x p_x \rho^x$, where p_x is the marginal probability of $X = x$ under distribution D . ρ is the density matrix of the average first message of \mathcal{P} under distribution D .

Stage 1. We construct a $[t; c; l_1, \dots, l_t]^A$ safe coinless quantum protocol \mathcal{P}' for f with average error under distribution D $\epsilon_D^{\mathcal{P}'} \leq \epsilon_D^{\mathcal{P}} + (1/2)((4 \ln 2)I(X : M))^{1/2}$, and where the density matrix of the first message is independent of the input x to Alice. We now describe the protocol \mathcal{P}' . Suppose Alice is given $x \in \mathcal{X}$ and Bob is given $y \in \mathcal{Y}$. The qubits of AM are initialised to zero. Alice applies a unitary transformation U'_x on AM in order to prepare a purification $|\phi^x\rangle$ of ρ . She then sends the qubits of M as her first message to Bob. After this, Alice and Bob behave as in protocol \mathcal{P} . Hence in \mathcal{P}' the density matrix of the first message is independent of Alice’s input.

Let us now compare the situation in protocols \mathcal{P} and \mathcal{P}' when Alice’s input is x , Bob’s input is y , Alice has finished preparing the state $|\phi^x\rangle$, but no communication has taken place as yet. In protocol \mathcal{P} , the state of AMB at this point in time is $|\psi^x\rangle_{AM}|\mathbf{0}\rangle_B$. In protocol \mathcal{P}' , the state of AMB at this point in time is $|\phi^x\rangle_{AM}|\mathbf{0}\rangle_B$. Hence, $\epsilon_{x,y}^{\mathcal{P}'} \leq \epsilon_{x,y}^{\mathcal{P}} + (1/2)\| |\psi^x\rangle\langle\psi^x| - |\phi^x\rangle\langle\phi^x| \|_1$. Let q_{xy} denote the probability that $(X, Y) = (x, y)$ under distribution D . Then, the average error of \mathcal{P}' under distribution D

$$\begin{aligned} \epsilon_D^{\mathcal{P}'} &= \sum_{x,y} q_{xy} \epsilon_{x,y}^{\mathcal{P}'} \leq \sum_{x,y} q_{xy} \left(\epsilon_{x,y}^{\mathcal{P}} + \frac{1}{2} \| |\psi^x\rangle\langle\psi^x| - |\phi^x\rangle\langle\phi^x| \|_1 \right) = \epsilon_D^{\mathcal{P}} + \frac{1}{2} \sum_x p_x \| |\psi^x\rangle\langle\psi^x| - |\phi^x\rangle\langle\phi^x| \|_1 \\ &\leq \epsilon_D^{\mathcal{P}} + \frac{1}{2} ((4 \ln 2)I(X : M))^{1/2}. \end{aligned}$$

The last inequality follows by taking $|\phi^x\rangle$ to be the purifications promised by Fact 9.

Stage 2A. We now construct a $[t + 1; c + l_1; 0, 0, l_2, \dots, l_t]^B$ safe coinless quantum protocol \mathcal{Q}' for f with $\epsilon_{x,y}^{\mathcal{Q}'} = \epsilon_{x,y}^{\mathcal{P}'}$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$. Suppose Alice is given $x \in \mathcal{X}$ and Bob is given $y \in \mathcal{Y}$. Let A_1 denote all the qubits of A , except the last $l_1 + c$ qubits. Let A_2 denote the last $l_1 + c$ qubits of A . Thus, $A = A_1 A_2$. In protocol \mathcal{Q}' , Alice initially starts with work qubits A_1 only, and Bob initially starts with work qubits BMA_2 . The work qubits of Alice and Bob are initialised to zero. Bob commences protocol \mathcal{Q}' by constructing a canonical purification $|\eta\rangle_{MA_2}$ of ρ , where the reduced density matrix of M is ρ . Bob then sends A_2 to Alice. The density matrix of A_2 is independent of the inputs x, y (in fact, if $|\eta\rangle_{MA_2}$ is the Schmidt purification then the reduced density matrix of A_2 is also ρ). After receiving A_2 , Alice applies a unitary transformation V_x on A so that the state vector of AM becomes $|\phi^x\rangle_{AM}$. The existence of such a V_x follows from Fact 8. The global state of Alice’s and Bob’s qubits at this point in protocol \mathcal{Q}'

is the same as the global state of Alice's and Bob's qubits at the point in protocol \mathcal{P}' just after Alice has sent her first message to Bob. Bob now treats M as if it were the first message of Alice in \mathcal{P}' , and proceeds to compute his response N (the qubits of N are a subset of the qubits of MB) of length l_2 . Bob sends N to Alice and after this protocol \mathcal{Q}' proceeds as protocol \mathcal{P}' . In \mathcal{Q}' Bob starts the communication, the communication goes on for $t + 1$ rounds, the first message of Bob of length $l_1 + c$ viz. A_2 is a safe message, and the first message of Alice is zero qubits long.

Stage 2B. We finally construct our desired $[t - 1; c + l_1; l_2, \dots, l_t]^B$ safe coinless quantum protocol \mathcal{Q} for f with $\epsilon_{x,y}^{\mathcal{Q}} = \epsilon_{x,y}^{\mathcal{Q}'}$, for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$. In protocol \mathcal{Q} Bob, after doing the same computations as in \mathcal{Q}' , first sends as a single message the $(l_1 + c) + l_2$ qubits A_2N . After receiving A_2N , Alice applies V_x on A followed by her appropriate unitary transformation on AN viz. the unitary transformation of Alice in \mathcal{Q}' on the qubits of AN after she has received the first two messages of Bob. The global state of all the qubits of Alice and Bob at this point in protocol \mathcal{Q} is the same as the global state of all the qubits of Alice and Bob at the point in protocol \mathcal{Q}' just after Alice has finished generating her second message but before she has sent it to Bob. This is because unitary transformations on disjoint sets of qubits commute. After this, protocol \mathcal{Q} proceeds as protocol \mathcal{Q}' . In protocol \mathcal{Q} Bob starts the communication, the communication goes on for $t - 1$ rounds, and the first message of Bob of length $(l_1 + c) + l_2$ viz. A_2N contains a safe overhead viz. A_2 of $l_1 + c$ qubits.

This completes the proof of Lemma 3. \square

Remark. The proof of Lemma 3 requires the global state of all the qubits of Alice and Bob to be pure at all times during the execution of protocol \mathcal{P} , when Alice's and Bob's inputs are in computational basis states. This is because we use the machinery of purifications in the proof. The proof also ensures that the purity property holds for the final protocol \mathcal{Q} .

4. The round elimination lemma

We now prove the classical and quantum versions of our strong round elimination lemma in Sections 4.1 and 4.2 respectively. The round elimination lemma is stated for public coin protocols only. Since a public coin classical randomised protocol can be converted to a private coin classical randomised protocol at the expense of an additive increase in the communication complexity by at most logarithm of the total bit size of the inputs [30], we also get a similar round elimination lemma for private coin classical randomised protocols. A similar statement about round elimination can be made for safe coinless quantum protocols; for such protocols the safe overhead increases by an additional additive term that is at most logarithm of the total bit size of the inputs. But since the statement of the round elimination lemma is cleanest for public coin protocols, we give it below for such protocols only.

4.1. Classical

Lemma 4 (Round elimination lemma, classical version). *Suppose $f: \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ is a function. Suppose the communication game $f^{(n),A}$ has a $[t; l_1, \dots, l_t]^A$ public coin randomised protocol with error less than δ . Then there is a $[t - 1; l_2, \dots, l_t]^B$ public coin randomised protocol for f with error less than $\epsilon \triangleq \delta + (1/2)(2l_1 \ln 2/n)^{1/2}$.*

Proof. Suppose the given protocol for $f^{(n),A}$ has error $\tilde{\delta} < \delta$. Define $\tilde{\epsilon} \triangleq \tilde{\delta} + (1/2)(2l_1 \ln 2/n)^{1/2}$. To prove the round elimination lemma it suffices to give, by the harder direction of Yao's minimax lemma [43], for any probability distribution D on $\mathcal{X} \times \mathcal{Y}$, a $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol \mathcal{P} for f with $\epsilon_D^{\mathcal{P}} \leq \tilde{\epsilon} < \epsilon$. To this end, we will first construct a probability distribution D^* on $\mathcal{X}^n \times [n] \times \mathcal{Y}$ as follows: Choose $i \in [n]$ uniformly at random. Choose independently, for each $j \in [n]$, $(x_j, y_j) \in \mathcal{X} \times \mathcal{Y}$ according to distribution D . Set $y = y_i$ and throw away $y_j, j \neq i$. By the easier direction of Yao's minimax lemma, we get a $[t; l_1, \dots, l_t]^A$ deterministic protocol \mathcal{P}^* for $f^{(n),A}$ with $\epsilon_{D^*}^{\mathcal{P}^*} \leq \tilde{\delta} < \delta$. In \mathcal{P}^* , Alice gets $x_1, \dots, x_n \in \mathcal{X}$, Bob gets $i \in [n]$, $y \in \mathcal{Y}$ and copies of x_1, \dots, x_{i-1} . We shall construct the desired protocol \mathcal{P} from the protocol \mathcal{P}^* .

In \mathcal{P}^* , let Alice's and Bob's inputs be distributed according to D^* . Let the input to Alice be denoted by the random variable $X \triangleq X_1 \cdots X_n$, where X_i is the random variable corresponding to the i th input to Alice. Let the random variables Y, \mathcal{I} correspond to the inputs y, i respectively of Bob. Let M denote the random variable corresponding to the first message of Alice in \mathcal{P}^* . Define probability distribution $D_{i;x_1, \dots, x_{i-1}}^*$ on $\mathcal{X}^n \times [n] \times \mathcal{Y}$ to be the distribution D^* conditioned on $\mathcal{I} = i$ and $X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1}$. Define probability distribution $D_{i;y;x_1, \dots, x_i}^*$ on $\mathcal{X}^n \times [n] \times \mathcal{Y}$

to be the distribution D^* conditioned on $\mathcal{I} = i$, $Y = y$ and $X_1, \dots, X_i = x_1, \dots, x_i$. Let $\epsilon_{D^*;i;x_1,\dots,x_{i-1}}^{\mathcal{P}^*}$ denote the average error of \mathcal{P}^* under distribution $D^*_{i;x_1,\dots,x_{i-1}}$. Using Fact 4 and the fact that under distribution D^* X_1, \dots, X_n are independent random variables, we get that

$$\mathbf{E}_{i,X} [I((X_i : M) \mid X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1})] = \mathbf{E}_i [I(X_i : M X_1 \cdots X_{i-1})] = \frac{I(X : M)}{n} \leq \frac{l_1}{n}. \quad (1)$$

Also,

$$\tilde{\delta} \geq \epsilon_{D^*}^{\mathcal{P}^*} = \mathbf{E}_{i,X} [\epsilon_{D^*;i;x_1,\dots,x_{i-1}}^{\mathcal{P}^*}]. \quad (2)$$

Above, the expectations are under distribution D^* and the mutual informations are for protocol \mathcal{P}^* with its inputs distributed according to D^* .

For any $i \in [n]$, $x_1, \dots, x_{i-1} \in \mathcal{X}$, let us now define the $[t; l_1, \dots, l_t]^A$ private coin randomised protocol $\mathcal{P}'_{i;x_1,\dots,x_{i-1}}$ for the function f in terms of protocol \mathcal{P}^* as follows: Alice is given $x \in \mathcal{X}$ and Bob is given $y \in \mathcal{Y}$. Bob sets $\mathcal{I} = i$, and both Alice and Bob set $X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1}$. Alice tosses a fresh private coin to choose $X_{i+1}, \dots, X_n \in \mathcal{X}$, where each X_j , $i + 1 \leq j \leq n$, is chosen independently according to the marginal distribution on \mathcal{X} induced by D . Alice sets $X_i = x$ and Bob sets $Y = y$. They then run protocol \mathcal{P}^* on these inputs. The probability that $\mathcal{P}'_{i;x_1,\dots,x_{i-1}}$ makes an error for an input (x, y) , $\epsilon_{x,y}^{\mathcal{P}'_{i;x_1,\dots,x_{i-1}}}$, is the average probability of error of \mathcal{P}^* under distribution $D^*_{i;y;x_1,\dots,x_i}$. Hence, the average probability of error of $\mathcal{P}'_{i;x_1,\dots,x_{i-1}}$ under distribution D ,

$$\epsilon_D^{\mathcal{P}'_{i;x_1,\dots,x_{i-1}}} = \epsilon_{D^*;i;x_1,\dots,x_{i-1}}^{\mathcal{P}^*}. \quad (3)$$

Let M' denote the random variable corresponding to Alice's first message and X' denote the random variable X_i , when Alice's and Bob's inputs are distributed according to D in $\mathcal{P}'_{i;x_1,\dots,x_{i-1}}$. Then

$$I((X_i : M) \mid X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1}) = I(X' : M'), \quad (4)$$

where the left-hand side refers to the mutual information in protocol \mathcal{P}^* when its inputs are distributed according to D^* and the right-hand side refers to the mutual information in protocol $\mathcal{P}'_{i;x_1,\dots,x_{i-1}}$ when its inputs are distributed according to D .

Using Lemma 2 and Eqs. (3) and (4), we get a $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol $\mathcal{P}_{i;x_1,\dots,x_{i-1}}$ for f with

$$\begin{aligned} \epsilon_D^{\mathcal{P}_{i;x_1,\dots,x_{i-1}}} &\leq \epsilon_D^{\mathcal{P}'_{i;x_1,\dots,x_{i-1}}} + \frac{1}{2}((2 \ln 2)I(X' : M'))^{1/2} \\ &= \epsilon_{D^*;i;x_1,\dots,x_{i-1}}^{\mathcal{P}^*} + \frac{1}{2}((2 \ln 2)I((X_i : M) \mid X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1}))^{1/2}. \end{aligned} \quad (5)$$

We have that (note that the expectations below are under distribution D^* and the mutual informations are for protocol \mathcal{P}^* with its inputs distributed according to D^*)

$$\begin{aligned} \mathbf{E}_{i,X} [\epsilon_D^{\mathcal{P}_{i;x_1,\dots,x_{i-1}}}] &\leq \mathbf{E}_{i,X} [\epsilon_{D^*;i;x_1,\dots,x_{i-1}}^{\mathcal{P}^*}] + \frac{1}{2} \mathbf{E}_{i,X} [((2 \ln 2)I((X_i : M) \mid X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1}))^{1/2}] \\ &\leq \mathbf{E}_{i,X} [\epsilon_{D^*;i;x_1,\dots,x_{i-1}}^{\mathcal{P}^*}] + \frac{1}{2} \left((2 \ln 2) \mathbf{E}_{i,X} [I((X_i : M) \mid X_1, \dots, X_{i-1} = x_1, \dots, x_{i-1})] \right)^{1/2} \\ &\leq \tilde{\delta} + \frac{1}{2} \left(\frac{2l_1 \ln 2}{n} \right)^{1/2} = \tilde{\epsilon}. \end{aligned}$$

The first inequality follows from (5), the second inequality follows from the concavity of the square root function and the last inequality from (1) and (2).

Thus, we can immediately see that there exist $i \in [n]$ and $x_1, \dots, x_{i-1} \in \mathcal{X}$ such that $\epsilon_D^{\mathcal{P}_{i;x_1,\dots,x_{i-1}}} \leq \tilde{\epsilon}$. Let $\mathcal{P} \triangleq \mathcal{P}_{i;x_1,\dots,x_{i-1}}$. \mathcal{P} is our desired $[t - 1; l_2, \dots, l_t]^B$ deterministic protocol for f with $\epsilon_D^{\mathcal{P}} \leq \tilde{\epsilon}$, thus completing the proof of the round elimination lemma. \square

4.2. Quantum

Lemma 5 (Round elimination lemma, quantum version). *Suppose $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$ is a function. Suppose the communication game $f^{(n),A}$ has a $[t; c; l_1, \dots, l_t]^A$ safe public coin quantum protocol with error less than δ . Then there is a $[t-1; c+l_1; l_2, \dots, l_t]^B$ safe public coin quantum protocol for f with error less than $\epsilon \triangleq \delta + (1/2)(8l_1 \ln 2/n)^{1/2}$.*

Proof (Sketch). The proof is very similar to the proof of Lemma 4. We just point out some important things below. Note that in the quantum setting, the upper bound in (1) is $\frac{2l_1}{n}$ by Fact 6(c). Also, in the definition of protocol $\mathcal{P}'_{i;x_1, \dots, x_{i-1}}$, instead of feeding probabilistic mixtures for the inputs X_{i+1}, \dots, X_n , Alice feeds appropriate pure states (pure states that would give the correct probabilistic mixture were they to be measured in the computational basis). Equations 3 and 4 continue to hold because protocol \mathcal{P}^* is secure. Since the global state of all the qubits of Alice and Bob is pure at all times during the execution of protocol $\mathcal{P}'_{i;x_1, \dots, x_{i-1}}$ when Alice's and Bob's inputs are in computational basis states, Lemma 3 can now be used to get the $[t-1; c+l_1; l_2, \dots, l_t]^B$ safe coinless quantum protocol $\mathcal{P}_{i;x_1, \dots, x_{i-1}}$ for f . \square

5. Optimal lower bounds for predecessor

In this section, we prove our (optimal) lower bounds on the query complexity of static predecessor searching in the cell probe model with randomised or address-only quantum query schemes.

Theorem 1. *Suppose there is a $(n^{O(1)}, (\log m)^{O(1)}, t)$ randomised cell probe scheme for the (m, n) -static predecessor problem with error probability less than $1/3$. Then, there is a constant c such that:*

- (a) *for all sufficiently large m , a value of n can be chosen so that $t \geq \frac{c \log \log m}{\log \log \log m}$ as a function of m ;*
- (b) *for all sufficiently large n , a value of m can be chosen so that $t \geq \sqrt{\frac{c \log n}{\log \log n}}$ as a function of n .*

The same lower bound also holds for address-only quantum cell probe schemes for static predecessor searching.

Proof. The proof is similar to the proof of the lower bound for predecessor in [29], but with different parameters, and using our stronger round elimination lemma in its classical version (Lemma 4).

By Fact 1, it suffices to consider communication protocols for the rank parity communication game $\text{PAR}_{\log m, n}$. Let $n = 2^{(\log \log m)^2 / \log \log \log m}$. Let $c_1 \triangleq (2 \ln 2)6^2$. For any given constants $c_2, c_3 \geq 1$, define

$$a \triangleq c_2 \log n, \quad b \triangleq (\log m)^{c_3}, \quad t \triangleq \frac{\log \log m}{(c_1 + c_2 + c_3) \log \log \log m}.$$

We will show that $\text{PAR}_{\log m, n}$ does not have $(2t, a, b)^A$ public coin randomised communication protocols with error less than $1/3$, thus proving both the desired lower bounds for the predecessor problem.

Given a $(2t, a, b)^A$ public coin protocol for $\text{PAR}_{\log m, n}$ with error probability at most δ , we can get a $(2t, a, b)^A$ public coin protocol for $\text{PAR}_{\frac{c_1 a t^2}{\log m}, n}^{(c_1 a t^2), A}$ with error probability at most δ by Fact 2. Using Lemma 4, we get a $(2t-1, a, b)^B$ public coin protocol for $\text{PAR}_{\frac{\log m}{c_1 a t^2}, n}^{\log m}$, but the error probability increases to at most $\delta + (12t)^{-1}$. By Fact 3, we get a $(2t-1, a, b)^B$ public coin protocol for $\text{PAR}_{\frac{\log m}{c_1 a t^2}, \frac{n}{c_1 b t^2}}^{(c_1 b t^2), B}$ with error probability at most $\delta + (12t)^{-1}$. From the given values of the parameters, we see that $\frac{\log m}{(2c_1 a t^2)^t} \geq \log(c_1 b t^2) + 1$. This implies that we also have a $(2t-1, a, b)^B$ public coin protocol for $\text{PAR}_{\frac{\log m}{2c_1 a t^2}, \frac{n}{c_1 b t^2}}^{(c_1 b t^2), B}$ with error probability at most $\delta + (12t)^{-1}$. Using Lemma 4 again, we get a $(2t-2, a, b)^A$ public coin protocol for $\text{PAR}_{\frac{\log m}{2c_1 a t^2}, \frac{n}{c_1 b t^2}}$, but the error probability increases to at most $\delta + 2(12t)^{-1}$.

We do the above steps repeatedly. We start off with a $(2t, a, b)^A$ public coin protocol for $\text{PAR}_{\log m, n}$ with error probability less than $1/3$. After applying the above steps i times, we get a $(2t - 2i, a, b)^A$ public coin protocol for $\text{PAR}_{\frac{\log m}{(2c_1 a t^2)^i}, \frac{n}{(c_1 b t^2)^i}}$ with error probability less than $1/3 + 2i(12t)^{-1}$.

By applying the above steps t times, we finally get a $(0, a, b)^A$ public coin randomised protocol for the problem $\text{PAR}_{\frac{\log m}{(2c_1 a t^2)^t}, \frac{n}{(c_1 b t^2)^t}}$ with error probability less than $1/3 + 2t(12t)^{-1} = 1/2$. From the given values of the parameters, we see that $\frac{\log m}{(2c_1 a t^2)^t} \geq (\log m)^{\Omega(1)}$ and $\frac{n}{(c_1 b t^2)^t} \geq n^{\Omega(1)}$. Thus, we get a zero round protocol for a rank parity problem on a non-trivial domain with error probability less than $1/2$, which is a contradiction.

In the above proof, we are tacitly ignoring “rounding off” problems. We remark that this does not affect the correctness of the proof.

Finally we observe that by using Lemma 3, one can prove that the same lower bound holds for address-only quantum cell probe schemes for static predecessor. \square

6. The ‘greater-than’ problem

We illustrate another application of the round elimination lemma to communication complexity by proving improved rounds versus communication tradeoffs for the ‘greater-than’ problem.

Theorem 2. *The bounded error public coin randomised t -round communication complexity of GT_n is lower bounded by $\Omega(n^{1/t}t^{-2})$. For bounded error quantum protocols with input-independent prior entanglement for GT_n , we have a lower bound of $\Omega(n^{1/t}t^{-1})$.*

Proof. We recall the following reduction from $\text{GT}_{n/k}^{(k),A}$ to GT_n (see [29]): In $\text{GT}_{n/k}^{(k),A}$, Alice is given $x_1, \dots, x_k \in \{0, 1\}^{n/k}$, Bob is given $i \in [k]$, $y \in \{0, 1\}^{n/k}$, and copies of x_1, \dots, x_{i-1} , and they have to communicate and decide if $x_i > y$. To reduce $\text{GT}_{n/k}^{(k),A}$ to GT_n , Alice constructs $\hat{x} \in \{0, 1\}^n$ by concatenating x_1, \dots, x_k , Bob constructs $\hat{y} \in \{0, 1\}^n$ by concatenating $x_1, \dots, x_{i-1}, y, 1^{n(1-i/k)}$. It is easy to see that $\hat{x} > \hat{y}$ iff $x_i > y$.

Suppose there is a t -round bounded error public coin randomised protocol for GT_n with communication complexity c . We can think of the protocol as a $(t, c, c)^A$ public coin randomised protocol with error probability less than $1/3$. Suppose $n \geq k^t$, where $k \triangleq (2 \ln 2)(3t)^2 c$. Applying the self-reduction and Lemma 4 alternately for t stages gives us a zero round protocol for the ‘greater-than’ problem on a non-trivial domain with error probability less than $1/2$, which is a contradiction. Thus, $c = \Omega(n^{1/t}t^{-2})$.

In the above proof, we are tacitly ignoring “rounding off” problems. We remark that this does not affect the correctness of the proof.

Finally, we observe that for bounded error quantum protocols with input-independent prior entanglement for GT_n , one can improve the lower bound to $\Omega(n^{1/t}t^{-1})$ by exploiting the fact that by definition, a quantum protocol sends fixed length messages independent of the input. \square

7. Discussion

In this paper, we have proved a lower bound for the randomised query complexity and address-only quantum query complexity of a cell probe scheme for the static predecessor searching problem. Our lower bound matches the deterministic cell probe upper bound of Beame and Fich when we seek bounds that are purely a function of the universe size or the subset size, but not both. We derived our lower bound by proving a strong round elimination lemma in communication complexity. Our round elimination lemma improves on the round elimination lemma of Miltersen, Nisan, Safra and Wigderson [29], and is crucial to proving our optimal lower bound for predecessor searching. Our strong round elimination lemma also gives us improved rounds versus communication tradeoffs for the ‘greater-than’ problem. We believe that our round elimination lemma is of independent interest and should have further applications. In fact, it has already been used in subsequent work on cell probe lower bounds for various data structure problems, for example, as part of Chakrabarti and Regev’s [6] lower bound for randomised query schemes for the approximate nearest neighbour problem in the Hamming cube. Our paper does not manage to show optimal lower bounds on the number of queries for the predecessor problem for arbitrary values of the universe size, subset size, number of

cells and word size. This deficiency has been remedied by Pătraşcu and Thorup [33,34] in very recent work, who show optimal lower bounds for randomised query schemes for static predecessor for arbitrary values of the problem parameters, using a novel direct sum based approach.

We believe that our work brings out an interesting fact. Sometimes, in order to prove lower bound results, it helps to work in a more general model of computation. Of course, this makes the task of proving lower bounds harder, but also we now have more tools and techniques at our disposal. This sometimes enables us to attack the problem in a clearer fashion, without letting irrelevant details about the restricted model distract us. In our case, our attempt to prove a lower bound result in the more general address-only quantum cell probe model led us to make better use of powerful tools from information theory, which finally enabled us to prove optimal lower bounds for predecessor in the randomised cell probe model! Also, the information-theoretic approach gives us a simpler and clearer lower bound proof as compared to previous lower bound proofs for predecessor. By now, several instances have come to light where new classical lower bounds have been proved by quantum arguments. The most spectacular example so far has been the exponential lower bound for general two-query locally decodable classical codes by Kerenidis and de Wolf [20]. Their lower bound proof uses quantum arguments and it is not clear how to dequantise it. Also, no other proof, classical or quantum, of this result is known!

The lower bound for predecessor searching for quantum cell probe schemes works only if the query scheme is *address-only*. If the query scheme is not address-only the asymmetry in the message lengths of Alice and Bob in the corresponding quantum communication protocol breaks down, and it becomes impossible to prove query lower bounds via communication complexity. Proving non-trivial lower bounds for static data structure problems for general quantum query schemes is thus an important open problem. Another open problem would be to extend the optimal lower bounds of Pătraşcu and Thorup [33,34] for the predecessor problem to address-only quantum query schemes.

Acknowledgments

We thank Rahul Jain, Hartmut Klauck and Peter Bro Miltersen for helpful discussions and feedback, Amit Chakrabarti for sending us an early copy of [6], Mihai Pătraşcu for sending us preliminary versions of [33,34], Ashwin Nayak for helpful discussions and for pointing out the reference [27], and Jaikumar Radhakrishnan for reading an early draft of [36] and for enlightening discussions. We also thank the anonymous referee for helpful feedback and pointing out that our lower bounds are optimal only when viewed purely as functions of universe size or subset size, but not both.

Appendix A. The average encoding theorem

A.1. Classical

In this subsection, we give a self-contained classical proof, without using quantum information theory, of Fact 5. We first recall the definition of a classical information-theoretic quantity called relative entropy, also known as *Kullback–Leibler divergence*.

Definition 10 (*Relative entropy*). Let P and Q be probability distributions on the same sample space Ω . The *relative entropy* between P and Q is defined as

$$S(P \parallel Q) \triangleq \sum_{x \in \Omega} P(x) \log \left(\frac{P(x)}{Q(x)} \right).$$

We now require a non-trivial fact from classical information theory, which upper bounds the total variation distance of a pair of probability distributions in terms of their relative entropy. A proof of the fact can be found in [8, Lemma 12.6.1].

Fact 10. Let P and Q be probability distributions on the same finite sample space Ω . Then,

$$\|P - Q\|_1 \leq \sqrt{(2 \ln 2) S(P \parallel Q)}.$$

We can now prove Fact 5.

Fact 5 (Average encoding theorem, classical version). Let X, M be correlated random variables. Let p_x denote the (marginal) probability that $X = x$, and Π^x denote the conditional distribution of M given that $X = x$. Let Π denote the (marginal) probability distribution of M , i.e. $\Pi = \sum_x p_x \Pi^x$. Then,

$$\sum_x p_x \|\Pi^x - \Pi\|_1 \leq \sqrt{(2 \ln 2) I(X : M)}.$$

Proof. Let \mathcal{X}, \mathcal{M} be the finite ranges of random variables X, M respectively. We define two probability distributions P, Q on $\mathcal{X} \times \mathcal{M}$. In distribution P , the probability of $(x, m) \in \mathcal{X} \times \mathcal{M}$ is $p_x \cdot \pi_m^x$, where π_m^x is the conditional probability that $M = m$ given that $X = x$. In distribution Q , the probability of $(x, m) \in \mathcal{X} \times \mathcal{M}$ is $p_x \cdot \pi_m$, where π_m is the (marginal) probability that $M = m$, i.e. $\pi_m = \sum_x p_x \pi_m^x$.

It is easy to check that $S(P \parallel Q) = I(X : M)$ and $\|P - Q\|_1 = \sum_x p_x \|\Pi_x - \Pi\|_1$. The result now follows by applying Fact 10 to P and Q . \square

A.2. Quantum

In this subsection, we give a proof of Fact 9. We first recall some basic definitions and facts from quantum information theory. Let ρ and σ be density matrices in the same finite-dimensional Hilbert space \mathcal{H} . The *fidelity* (also called Uhlmann’s transition probability or the Bhattacharya coefficient) of ρ and σ is defined as $B(\rho, \sigma) \triangleq \|\sqrt{\rho} \sqrt{\sigma}\|_1$. The *von Neumann relative entropy* between ρ and σ is defined as $S(\rho \parallel \sigma) \triangleq \text{Tr}(\rho(\log \rho - \log \sigma))$.

Jozsa [19] gave an elementary proof for finite-dimensional Hilbert spaces of the following basic and remarkable property about fidelity.

Fact 11. Let ρ, σ be density matrices in the same Hilbert space \mathcal{H} . Let \mathcal{K} be a disjoint Hilbert space such that $\dim(\mathcal{K}) \geq \dim(\mathcal{H})$. Then for any purification $|\psi\rangle$ of ρ in $\mathcal{H} \otimes \mathcal{K}$, there exists a purification $|\phi\rangle$ of σ in $\mathcal{H} \otimes \mathcal{K}$ such that $B(\rho, \sigma) = |\langle \psi | \phi \rangle|$.

For two probability distributions P, Q on the same sample space Ω , their fidelity is defined as $B(P, Q) \triangleq \sum_{x \in \Omega} \sqrt{P(x)Q(x)}$. We will need the following result about fidelity proved by Fuchs and Caves [13].

Fact 12. Let ρ, σ be density matrices in the same Hilbert space \mathcal{H} . Then $B(\rho, \sigma) = \inf_{\mathcal{M}} B(\mathcal{M}\rho, \mathcal{M}\sigma)$, where \mathcal{M} ranges over POVMs on \mathcal{H} . In fact, the infimum above can be attained by a complete von Neumann measurement on \mathcal{H} .

The following fundamental fact (see e.g. [31]) states that the relative entropy can only decrease on performing a measurement.

Fact 13 (Monotonicity of relative entropy). Let ρ, σ be density matrices in the same finite-dimensional Hilbert space \mathcal{H} . Let \mathcal{M} be a POVM on \mathcal{H} . Then, $S(\mathcal{M}\rho \parallel \mathcal{M}\sigma) \leq S(\rho \parallel \sigma)$.

We require the following explicit expression for the trace distance of two pure states (see e.g. [31]).

Fact 14. For pure states $|\psi\rangle$ and $|\phi\rangle$ in the same Hilbert space, $\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1 = 2\sqrt{1 - |\langle\psi|\phi\rangle|^2}$.

The following information-theoretic fact follows easily from the definitions.

Fact 15. Let X, M be quantum systems with some joint density matrix, where X is a classical random variable. Let p_x denote the (marginal) probability that $X = x$ and ρ^x denote the conditional density matrix of M given that $X = x$. Let ρ denote the reduced density matrix of M . Note that $\rho = \sum_x p_x \rho^x$. Then, $I(X : M) = \sum_x p_x S(\rho^x \parallel \rho)$.

We now recall the following direct connection between relative entropy and fidelity observed in the classical setting by [9], and in the quantum setting by Klauck (private communication).

Fact 16. Let ρ and σ be two density matrices in the same Hilbert space. Then,

$$1 - B(\rho, \sigma) \leq \frac{(\ln 2)S(\rho \parallel \sigma)}{2}.$$

Proof. Let \mathcal{M} be the complete von Neumann measurement that achieves the infimum in Fact 12. Let Ω denote the set of possible (classical) outcomes of \mathcal{M} . Define probability distributions $P \triangleq \mathcal{M}\rho$ and $Q \triangleq \mathcal{M}\sigma$. From Fact 13 and concavity of the log function it follows that

$$-\frac{S(\rho \parallel \sigma)}{2} \leq -\frac{S(P \parallel Q)}{2} = \sum_{x \in \Omega} P(x) \log \sqrt{\frac{Q(x)}{P(x)}} \leq \log \sum_{x \in \Omega} \sqrt{Q(x)P(x)} = \log B(P, Q) = \log B(\rho, \sigma).$$

Thus, $B(\rho, \sigma) \geq 2^{-S(\rho \parallel \sigma)/2} = \exp(-(\ln 2)S(\rho \parallel \sigma)/2) \geq 1 - ((\ln 2)S(\rho \parallel \sigma)/2)$. \square

We can now prove Fact 9.

Fact 9 (Average encoding theorem, quantum version). Let X, M be quantum systems with some joint density matrix. Let X be a classical random variable and p_x denote the (marginal) probability that $X = x$. Let \mathcal{H} denote the Hilbert space of M and ρ^x denote the conditional density matrix of M given that $X = x$. Let ρ denote the reduced density matrix of M . Note that $\rho = \sum_x p_x \rho^x$. Let \mathcal{K} denote a disjoint Hilbert space such that $\dim(\mathcal{K}) \geq \dim(\mathcal{H})$. Let $|\psi^x\rangle$ be purifications of ρ^x in $\mathcal{H} \otimes \mathcal{K}$. Then there exist purifications $|\phi^x\rangle$ of ρ in $\mathcal{H} \otimes \mathcal{K}$ such that

$$\sum_x p_x \left\| |\psi^x\rangle\langle\psi^x| - |\phi^x\rangle\langle\phi^x| \right\|_1 \leq \sqrt{(4 \ln 2)I(X : M)}.$$

Proof. Using the concavity of the square root function, Facts 15, 11, 14, 16, and the fact that fidelity is always at most 1, we get

$$\begin{aligned} \sqrt{I(X : M)} &= \sqrt{\sum_x p_x S(\rho^x \parallel \rho)} \geq \sum_x p_x \sqrt{S(\rho^x \parallel \rho)} \geq \sum_x p_x \sqrt{\frac{2(1 - B(\rho^x, \rho))}{\ln 2}} \\ &\geq \sum_x p_x \sqrt{\frac{1 - (B(\rho^x, \rho))^2}{\ln 2}} = \sum_x p_x \sqrt{\frac{1 - |\langle\psi^x|\phi^x\rangle|^2}{\ln 2}} = \sum_x p_x \frac{\| |\psi^x\rangle\langle\psi^x| - |\phi^x\rangle\langle\phi^x| \|_1}{\sqrt{4 \ln 2}}. \quad \square \end{aligned}$$

References

- [1] D. Aharonov, A. Kitaev, N. Nisan, Quantum circuits with mixed states, in: Proceedings of the 30th Annual ACM Symposium on Theory of Computing, 1998, pp. 20–30.
- [2] M. Ajtai, A lower bound for finding predecessors in Yao's cell probe model, *Combinatorica* 8 (3) (1988) 235–247.
- [3] A. Andersson, Faster deterministic sorting and searching in linear space, in: Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science, 1996, pp. 135–141.
- [4] P. Beame, F. Fich, Optimal bounds for the predecessor problem and related problems, *J. Comput. System Sci.* 65 (1) (2002) 38–72.
- [5] H. Buhrman, P.B. Miltersen, J. Radhakrishnan, S. Venkatesh, Are bitvectors optimal?, *SIAM J. Comput.* 31 (6) (2002) 1723–1744.
- [6] A. Chakrabarti, O. Regev, An optimal randomised cell probe lower bound for approximate nearest neighbour searching, in: Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, 2004, pp. 473–482.
- [7] R. Cleve, Wim. van Dam, M. Nielsen, A. Tapp, Quantum entanglement and the communication complexity of the inner product function, in: Proceedings of the 1st NASA International Conference on Quantum Computing and Quantum Communications, in: Lecture Notes in Comput. Sci., vol. 1509, Springer-Verlag, 1998, pp. 61–74.
- [8] T. Cover, J. Thomas, Elements of Information Theory, Wiley Series in Telecommunications, John Wiley and Sons, 1991.
- [9] D. Dacunha-Castelle, H. Heyer, B. Roynette, Ecole d'Eté de Probabilités de Saint-Flour VII, Lecture Notes in Math., vol. 678, Springer-Verlag, 1978.
- [10] E. Farhi, J. Goldstone, S. Gutmann, M. Sipser, Invariant quantum algorithms for insertion into an ordered list, manuscript at <http://arxiv.org/abs/quant-ph/9901059>, January 1999.
- [11] M. Fredman, J. Komlós, E. Szemerédi, Storing a sparse table with $O(1)$ worst case access time, *J. Assoc. Comput. Mach.* 31 (3) (1984) 538–544.
- [12] M. Fredman, D. Willard, Surpassing the information theoretic bound with fusion trees, *J. Comput. System Sci.* 47 (1993) 424–436.

- [13] C. Fuchs, C. Caves, Mathematical techniques for quantum communication theory, *Open Syst. Inf. Dyn.* 3 (3) (1995) 345–356; preliminary version at <http://arxiv.org/abs/quant-ph/9604001>.
- [14] L. Grover, A fast quantum mechanical algorithm for database search, in: *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, 1996, pp. 212–219.
- [15] P. Høyer, J. Neerbek, Y. Shi, Quantum complexities of ordered searching, sorting, and element distinctness, *Algorithmica* 34 (4) (2002) 429–448.
- [16] P. Indyk, R. Motwani, Approximate nearest neighbors: Towards removing the curse of dimensionality, in: *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, 1998, pp. 604–613.
- [17] R. Jain, J. Radhakrishnan, P. Sen, A direct sum theorem in communication complexity via message compression, in: *Proceedings of the 30th International Colloquium on Automata, Languages and Programming*, in: *Lecture Notes in Comput. Sci.*, vol. 2719, Springer-Verlag, 2003, pp. 300–315.
- [18] R. Jain, J. Radhakrishnan, P. Sen, Prior entanglement, message compression and privacy in quantum communication, in: *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, 2005, pp. 285–296.
- [19] R. Jozsa, Fidelity for mixed quantum states, *J. Modern Opt.* 41 (12) (1994) 2315–2323.
- [20] I. Kerenidis, R. de Wolf, Exponential lower bound for 2-query locally decodable codes via a quantum argument, *J. Comput. System Sci.* 69 (3) (2004) 395–420.
- [21] H. Klauck, Quantum communication complexity, in: *Proceedings of the Satellite Workshops at the 27th International Colloquium on Automata, Languages and Programming, Workshop on Boolean Functions and Applications (invited lecture)*, Carleton Scientific, Waterloo, Ontario, Canada, 2000, pp. 241–252; also available at <http://arxiv.org/abs/quant-ph/0005032>.
- [22] H. Klauck, A. Nayak, Amnon Ta-Shma, D. Zuckerman, Interaction in quantum communication and the complexity of set disjointness, in: *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, 2001, pp. 124–133.
- [23] I. Kremer, Quantum communication, Master's thesis, Hebrew University, 1995.
- [24] E. Kushilevitz, N. Nisan, *Communication Complexity*, Cambridge University Press, 1997.
- [25] E. Kushilevitz, R. Ostrovsky, Y. Rabani, Efficient search for approximate nearest neighbor in high-dimensional spaces, *SIAM J. Comput.* 30 (2) (2000) 457–474.
- [26] D. Liu, A strong lower bound for approximate nearest neighbor searching, *Inform. Process. Lett.* 92 (1) (2004) 23–29.
- [27] E. Maneva, Interactive communication on noisy channels, Bachelor's thesis, California Institute of Technology, 2001; available at <http://www.cs.berkeley.edu/~elitza/resume.html>.
- [28] P.B. Miltersen, Lower bounds for union-split-find related problems on random access machines, in: *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, 1994, pp. 625–634.
- [29] P. Bro Miltersen, Noam Nisan, S. Safra, A. Wigderson, On data structures and asymmetric communication complexity, *J. Comput. System Sci.* 57 (1) (1998) 37–49.
- [30] I. Newman, Private vs. common random bits in communication complexity, *Inform. Process. Lett.* 39 (1991) 67–71.
- [31] M. Nielsen, I. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2000.
- [32] N. Nisan, The communication complexity of threshold gates, in: *Combinatorics, Paul Erdős is Eighty*, vol. 1, János Bolyai Mathematical Society, Budapest, Hungary, 1993, pp. 301–315.
- [33] M. Pătraşcu, M. Thorup, Lower bounds for randomized predecessor search, unpublished manuscript, 2006.
- [34] M. Pătraşcu, M. Thorup, Time-space trade-offs for predecessor search, in: *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006, pp. 232–240.
- [35] P. Sen, Lower bounds for predecessor searching in the cell probe model, in: *Proceedings of the 18th Annual IEEE Conference on Computational Complexity*, 2003, pp. 73–83.
- [36] P. Sen, S. Venkatesh, Lower bounds in the quantum cell probe model, in: *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, in: *Lecture Notes in Comput. Sci.*, vol. 2076, Springer-Verlag, 2001, pp. 358–369.
- [37] D. Smirnov, Shannon's information methods for lower bounds for probabilistic communication complexity, Master's thesis, Moscow University, 1988.
- [38] P. van Emde Boas, Preserving order in a forest in less than logarithmic time and linear space, *Inform. Process. Lett.* 6 (1977) 80–82.
- [39] P. van Emde Boas, R. Kaas, E. Zijlstra, Design and implementation of an efficient priority queue, *Math. Systems Theory* 10 (1977) 99–127.
- [40] D. Willard, Log-logarithmic worst case range queries are possible in space $\theta(n)$, *Inform. Process. Lett.* 17 (1985) 81–84.
- [41] B. Xiao, New bounds in cell probe model, PhD thesis, University of California at San Diego, 1992.
- [42] A. Yao, Quantum circuit complexity, in: *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science*, 1993, pp. 352–361.
- [43] A.C.-C. Yao, Probabilistic computations: Towards a unified measure of complexity, in: *Proceedings of the 18th Annual IEEE Symposium on Foundations of Computer Science*, 1977, pp. 222–227.
- [44] A.C.-C. Yao, Should tables be sorted?, *J. Assoc. Comput. Mach.* 28 (3) (1981) 615–628.