A quantum information cost trade-off for the Augmented Index function[1]

Rahul Jain
National U. Singapore and C.Q.T.

Ashwin Nayak
U. Waterloo and Perimeter

Augmented Index is a variant of a basic problem in communication complexity, the Index function problem. In this variant, the player holding the index also receives a portion of the other party's input. More formally, one party, Alice, has an n-bit string x, and the other party, Bob, has an integer $k \in [n]$, the prefix $x[1, k-1]$ of x, and a bit $b \in \{0, 1\}$. Their goal is to compute the function $f_n(x, (k, x[1, k-1], b)) = x_k \oplus b$, i.e., to determine whether $b = x_k$ or not. This problem was studied in the one-way communication model as "serial encoding" [2, 24], and as "Augmented Index" [9, 14] and "Mountain problem" [21] in later works.

Communication problems involving the Index function and its variants capture a number of phenomena in the theory of computing, both classical and quantum, in addition to playing a fundamental role in the area of communication complexity [19]. For instance, they have been used to analyze data structures [22], the size of finite automata [3] and formulae [17], the length of locally decodable codes [15], learnability of quantum states [1], and sketching complexity [4]. Recently, phenomena in quantum information have been discovered via the Index function problem, e.g., information causality [27], a connection between non-locality and the uncertainty principle [26] and quantum ignorance [29].

Imagine that the two remotely situated parties, Alice and Bob, wish to compute the Augmented Index function by communicating with each other. Imagine further that they wish to keep as much of their respective inputs hidden from the other party as possible. This is precisely the kind of scenario that arises in the classic two millionaires' problem, in which Alice and Bob try to determine who is richer, without revealing more about their assets, or in Private Information Retrieval, in which a stock broker would like to look up the availability of a particular stock, without divulging which one it is. Intuitively, for any function that depends non-trivially on both sets of inputs, the two parties would perforce reveal some information about their inputs when no restrictions are placed on their computational abilities. In this work we establish a trade-off between the amount of (classical and) quantum information the two parties necessarily reveal about their inputs in the process of the computing Augmented Index. A surprising feature of this trade-off is that it holds even under a distribution on inputs on which the function value is *known in advance*. In fact, this is the price paid by any protocol that works correctly on a "hard" distribution. We show that in quantum protocols that compute Augmented Index correctly with constant error on the uniform distribution, either Alice reveals $\Omega(n/t)$ information about her n-bit input x, or Bob reveals $\Omega(1/t)$ information about his $(\log n)$-bit input k, where t is the number of messages in the protocol, even when the inputs are drawn from an "easy" distribution, the uniform distribution over inputs which evaluate to 0. In more detail, we start by defining appropriate notions of quantum information cost $(\mathrm{QIC}^{\mathsf{A}}_{\lambda}(\Pi), \mathrm{QIC}^{\mathsf{B}}_{\lambda}(\Pi))$ of a protocol Π for Augmented Index for the two players Alice (A) and Bob (B) with respect to the distribution λ, and then show:

Theorem 1 *In any two-party quantum communication protocol Π (with read-only behaviour on inputs and no intermediate measurements) for the* Augmented Index *function f_n that has t messages and makes constant error at most $\varepsilon \in [0, 1/4)$ on the uniform distribution μ over inputs, either $\mathrm{QIC}^{\mathsf{A}}_{\mu_0}(\Pi) \in \Omega(n/t)$ or $\mathrm{QIC}^{\mathsf{B}}_{\mu_0}(\Pi) \in \Omega(1/t)$, where μ_0 is the uniform distribution over $f_n^{-1}(0)$.*

Klauck [16] studied privacy in communication protocols in the setting of "honest but curious" players. The notion of privacy he considers is weaker than that in this work, as the privacy loss

[1] Full article available as ECCC Technical Report TR10-071, `http://eccc.hpi-web.de/report/2010/071/` , version 3, July 26, 2011.

is measured with respect to a (hard) distribution over inputs (rather than a superposition over inputs). A notion of information cost for INDEX was studied previously by Jain, Radhakrishnan, and Sen [12] in the context of privacy in communication. This notion differs from the one we study in two crucial respects. First, it is defined in terms of the hard distribution for the problem (uniform over all inputs). Second, the hard distribution is a product distribution. The techniques they develop seem not to be directly relevant to the problem at hand, as we deal with an easy and non-product distribution. Among known results, perhaps the one closest to Theorem 1 is a lower bound on information loss in the computation of the two-bit AND function due to Jain, Radhakrishnan, and Sen [11]. Information loss for AND is also defined in terms of an "easy" distribution, and analyzed for quantum protocols that are guaranteed to work in the worst case. We elaborate on another connection with this result below.

We devise a novel method for analyzing the information cost of f_n to arrive at Theorem 1. The proof we present shows how the conceptually simple and familiar ideas such as *average encoding* and *local transitions* may be brought to bear on AUGMENTED INDEX. These primitives were originally developed to prove properties of quantum protocols beyond the reach of previously known techniques [18], and have since then been specialized and applied to classical protocols with tremendous success. In fact, we first prove an analogue of Theorem 1 for classical protocols, which gives a stronger (optimal up to constant factors) trade-off for classical information cost. Classical protocols that produce the correct output with constant probability more than $3/4$ with respect to the distribution μ are such that either Alice reveals $\Omega(n)$ information about x or Bob reveals $\Omega(1)$ information about k, even under the distribution μ_0. The full version of this article compares the classical result to previous work by Magniez, Mathieu, and Nayak [21], and independent and concurrent work by Chakrabarti, Cormode, Kondapally, and McGregor [7] in detail. However, we emphasize that the approaches taken in these two works do not directly generalize to quantum protocols. They are based on analyzing the input distribution conditioned on the message transcript; no suitable quantum analogue of this technique is known.

The quantum information cost trade-off involves a number of subtleties. Unlike the classical information cost, which has been studied extensively (see, e.g., [8, 28, 5, 13]), it is not *a priori* clear what the appropriate definition of quantum information cost is. For one, the no-cloning principle [25] prevents the two parties from keeping a copy of the messages so there is no obvious notion of the history of the protocol. A notion of a transcript that encapsulates the history of a quantum protocol is instead the sequence of the joint states after each message exchange. Second, we consider the information contained about a *superposition* of inputs corresponding to the distribution of interest. This information is in general more than the information contained about a distribution over inputs, and the resulting notion seems to be necessary for the proof of the information cost trade-off we present. A third subtlety arises from the manner in which the input is distributed among the two parties. Alice and Bob share $x[1,k]$ when the inputs are restricted to $f_n^{-1}(0)$. Therefore when the input registers are initialized with the corresponding superpositions, the two parties already begin with some information about each other's input. Unlike in the classical case, this enables Alice to get information about the index k. The effect of sharing the prefix is identical to that of measuring the first k qubits of Alice's input superposition in the computational basis. This results in states of varying amount of von Neumann entropy for different indices, which leaks information about the index k. To quantify the information leaked by the protocol, we therefore imagine that there is a single quantum register that carries the superposition corresponding to x, and that Bob has read-only access to the relevant portion of this register. The information cost is then measured with respect to this register.

The intuition behind the lower bound on quantum information cost is as follows. Starting from an input pair on which the function evaluates to 0, if the information cost of any one party is

low and we carefully change her input, the other party's share of the state does not change much. Assume for simplicity that Alice produces the output of the protocol. We show that even when we simultaneously change the inputs with both parties simultaneously, resulting in a 1-input of the function, the perturbation to Alice's final state is also correspondingly small. This implies that the two information costs cannot be small simultaneously. The effect of simultaneously switching the two parties' inputs is captured in the classical case by the Cut-and-Paste lemma, which does not apply to quantum protocols. In the final piece of the argument above we instead appeal to the Local Transition Theorem and a hybrid argument. These are applied on a message-by-message basis, *à la* Jain *et al.* [11], and lead to a dependence of the information cost trade-off on the number of messages in the protocol. We are not aware of quantum protocols that beat the classical information bounds. However the dependence of the trade-off in Theorem 1 on the number of messages t may be inherent as is the case with Set Disjointness [11].

The classical version of this result, a trade-off for classical information cost with either $\Omega(n)$ information revealed by Alice or $\Omega(1)$ by Bob, has implications for the space required by *streaming algorithms*. It implies that streaming algorithms for certain context free properties (captured by the language DYCK(2)) need space $\Omega(\sqrt{n}/T)$ on inputs of length n, when allowed T unidirectional *passes* (sequential scans) over the input. In the context of classical computation, streaming algorithms were motivated by the growing need to process massive input data, which cannot fit entirely in computer memory [23]. Random access to such input is prohibitive, so ideally we would like to process it with a single sequential scan. Furthermore, during the computation, we are compelled to use space that is much smaller than the length of the input. Thus streaming algorithms scan the input sequentially only once (or a few times), while processing each input symbol quickly using a small amount of space. The need for algorithms of a similar simple form becomes more acute in the context of quantum computation. Streaming algorithms with quantum memory are the algorithms of choice in the absence of of prototypes with a large enough number of qubits and with long enough coherence times. Indeed, this has fuelled the study of quantum finite automata, which are precisely streaming algorithms that use constant space and time, and later works on quantum streaming algorithms [20, 10, 6].

The connection between the AUGMENTED INDEX function f_n and streaming algorithms for DYCK(2) was charted by Magniez *et al.* [21]. They map a streaming algorithm for DYCK(2) that uses space s to a multi-party communication protocol in which the messages are each of the same length s, and then bound s from below for protocols resulting from one-pass algorithms. The communication bound is derived using the information cost approach, which reduces the task to bounding from below the information cost of AUGMENTED INDEX with respect to the easy distribution μ_0. The proof of the connection between streaming algorithms and protocols for AUGMENTED INDEX does not extend immediately to the quantum case due to the stronger notion of information cost we define. Theorem 1 would however have similar consequences, if a certain information inequality holds. Consider unitary operations $U_{i,x,k}, V_{i,x}$ on m qubits, where $i \in [\ell], k \in [n]$, and $x \in \{0,1\}^n$, where $U_{i,x,k}$ depends on $x[1,k]$ and i, and $V_{i,x}$ depends on x and i. Let $\vec{K}$ be uniformly distributed over $[n]^\ell$, $\vec{X}$ be initialized to a uniform superposition over $\{0,1\}^{n\ell}$, and M be a quantum state over m qubits, obtained by successively applying $U_{i,X_i,K_i}V_{i,X_i}$ to $|\bar{0}\rangle$, controlled by the ith string X_i in $\vec{X}$, the ith index K_i in $\vec{K}$ for each $i \in [\ell]$. Under the conjecture that $\mathrm{I}(\vec{K} : M\vec{X}) \leq m$, we obtain a space lower bound for quantum streaming algorithms analogous to the classical one. We leave this potential implication for quantum streaming algorithms to future work.

References

[1] Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A, Mathematical, Physical & Engineering Sciences*, 463(2088):3089–3114, 2007.

[2] Andris Ambainis, Ashwin Nayak, Amnon Ta-Shma, and Umesh Vazirani. Dense quantum coding and a lower bound for 1-way quantum automata. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 376–383. ACM Press, May 1–4, 1999.

[3] Andris Ambainis, Ashwin Nayak, Amnon Ta-Shma, and Umesh Vazirani. Dense quantum coding and quantum finite automata. *Journal of the ACM*, 49(4):1–16, July 2002.

[4] Ziv Bar-Yossef, T. S. Jayram, Robert Krauthgamer, and Ravi Kumar. The sketching complexity of pattern matching. In Klaus Jansen, Sanjeev Khanna, José D. P. Rolim, and Dana Ron, editors, *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2004) and 8th International Workshop on Randomization and Computation (RANDOM 2004)*, volume 3122 of *Lecture Notes in Computer Science*, pages 261–272. Springer, 2004.

[5] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004. Special issue on FOCS 2002.

[6] Robin Blume-Kohout, Sarah Croke, and Daniel Gottesman. Streaming universal distortion-free entanglement concentration. Technical Report arXiv:0910.5952, arXiv.org Preprint, `http://arxiv.org/abs/0910.5952`, October 30, 2009.

[7] Amit Chakrabarti, Ranganath Kondapally Graham Cormode, and Andrew McGregor. Information cost tradeoffs for Augmented Index and streaming language recognition. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 387–396, Washington, DC, USA, 2010. IEEE Computer Society.

[8] Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew C.-C. Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 270–278, 2001.

[9] Khanh Do Ba, Piotr Indyk, Eric Price, and David P. Woodruff. Lower bounds for sparse recovery. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 1190–1197, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.

[10] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf. Exponential separation for one-way quantum communication complexity, with applications to cryptography. *SIAM Journal on Computing*, 38(5):1695–1708, 2008.

[11] Rahul Jain, Jaikumar Radhakrishnan, and Pranab Sen. A lower bound for the bounded round quantum communication complexity of Set Disjointness. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 220–229. IEEE Computer Society Press, Los Alamitos, CA, USA, 2003.

[12] Rahul Jain, Jaikumar Radhakrishnan, and Pranab Sen. A property of quantum relative entropy with an application to privacy in quantum communication. *Journal of the ACM*, 56(6):1–32, 2009.

[13] T. S. Jayram, Ravi Kumar, and D. Sivakumar. Two applications of information complexity. In *Proceedings of the Thirty-Fifth annual ACM Symposium on Theory of Computing*, pages 673–682. ACM, 2003.

[14] Daniel M. Kane, Jelani Nelson, and David P. Woodruff. On the exact space complexity of sketching and streaming small norms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 1161–1178, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.

[15] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes. *Journal of Computer and System Sciences*, 69(3):395–420, 2004. Special issue for STOC 2003.

[16] Hartmut Klauck. Quantum and approximate privacy. *Theory of Computing Systems*, 37(1):221–246, 2004.

[17] Hartmut Klauck. One-way communication complexity and the Nečiporuk lower bound on formula size. *SIAM Journal on Computing*, 37(2):552–583, 2007.

[18] Hartmut Klauck, Ashwin Nayak, Amnon Ta-Shma, and David Zuckerman. Interaction in quantum communication. *IEEE Transactions on Information Theory*, 53(6):1970–1982, June 2007.

[19] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, Cambridge, UK, 1997.

[20] François Le Gall. Exponential separation of quantum and classical online space complexity. In *Proceedings of the Eighteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '06, pages 67–73, New York, NY, USA, 2006. ACM.

[21] Frédéric Magniez, Claire Mathieu, and Ashwin Nayak. Recognizing well-parenthesized expressions in the streaming model. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing*, pages 261–270, New York, NY, June 6–8 2010. ACM Press.

[22] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On data structures and asymmetric communication complexity. *Journal of Computer and System Sciences*, 57(1):37–49, 1998.

[23] S. Muthukrishnan. *Data Streams: Algorithms and Applications*, volume 1, number 2 of *Foundations and Trends in Theoretical Computer Science*. Now Publishers Inc., Hanover, MA, USA, 2005.

[24] Ashwin Nayak. Optimal lower bounds for quantum automata and random access codes. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 369–376. IEEE Computer Society Press, October 17–19, 1999.

[25] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, UK, 2000.

[26] Jonathan Oppenheim and Stephanie Wehner. The uncertainty principle determines the non-locality of quantum mechanics. *Science*, 330(6007):1072–1074, 2010.

[27] Marcin Pawowski, Tomasz Paterek, Dagomir Kaszlikowski, Valerio Scarani, Andreas Winter, and Marek Żukowski. Information causality as a physical principle. *Nature*, 461:1101–1104, 2009.

[28] Michael Saks and Xiaodong Sun. Space lower bounds for distance approximation in the data stream model. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 360–369. ACM, 2002.

[29] Thomas Vidick and Stephanie Wehner. Does ignorance of the whole imply ignorance of the parts? Large violations of noncontextuality in quantum theory. *Physical Review Letters*, 107(030402), 2011.