A quantum information trade-off for Augmented Index

Ashwin Nayak

Joint work with Dave Touchette
(Waterloo)

Augmented Index (AI_n)

$x = x_1 x_2 \ldots x_n$

$k,\ x[1, k\text{-}1],\ b$

Is $x_k = b$?

Variant of Index function

Alice has an n-bit string x

Bob has the prefix $x[1, k\text{-}1]$, and a bit b

Goal: Compute $x_k \oplus b$

(Augmented) Index function

Fundamental problem with a rich history

- communication complexity [KN'97]
- data structures [MNSW'98]
- private information retrieval [CKGS'98]
- learnability of states [KNR'95, A'07]
- finite automata [ANTV'99]
- formula size [K'07]
- locally decodable codes [KdW'03]
- sketching e.g., [BJKK'04]
- information causality [PPKSWZ'09]
- non-locality and uncertainty principle [OW'10]
- quantum ignorance [VW'11] and more!

Connection with streaming algorithms

Magniez, Mathieu, N. '10:

- For Dyck(2): is an expression in two types of parentheses is well-formed ?
    - ([ ] ( )) is well-formed
    - ([ )( ]) is not well-formed
- Motivation: what is the complexity of problems beyond recognizing regular languages, say of context-free languages ?
- Dyck(2) is a canonical CFL, used in practice: e.g., checking well-formedness of large XML file

Streaming algorithms for Dyck(2)

Magniez, Mathieu, N.'10:

- A single pass randomized algorithm that uses $O((n \log n)^{1/2})$ space, $O(\text{polylog } n)$ time/ symbol
- 2-pass algorithm, uses $O(\log^2 n)$ space, $O(\text{polylog } n)$ time/ symbol, second pass in reverse
- Space usage of one-pass algorithm is optimal, via an information cost trade-off for Augmented Index (two-round)

Chakrabarti, Cormode, Kondapalli, McGregor '10; Jain, N.'10:

- Space usage of unidirectional T-pass algorithm is $n^{1/2} / T$
- Again, through information cost trade-off for Augmented Index, for an arbitrary number of rounds

Classical information trade-offs for AI_n

rounds	error	Alice reveals	or Bob reveals	Ref.
two, Alice starts	$1/(n \log n)$	$\Omega(n)$	$\Omega(n \log n)$	MMN'10
any no.	constant	$\Omega(n)$	$\Omega(1)$	CCKM'10 JN'10
any no.	constant	$\Omega(n/2^m)$	$\Omega(m)$	CK'11

- trade-offs w.r.t. uniform distribution over 0-inputs
- Internal information cost for classical protocols

Augmented Index AI_n

$x = x_1 x_2 \ldots x_n$

$k, \; x[1, k\text{-}1], \; b$

Is $x_k = b$?

- Simple protocols: Alice sends x or Bob sends k, b
- Can interpolate between the two:
  - Bob sends the m leading bits of k
  - Alice sends the corresponding block of x of length $n / 2^m$

Streaming algorithms

···0 1 0 1 1 0 0 1 0 1 0 1 0 1 1 1 0 0 1 0···

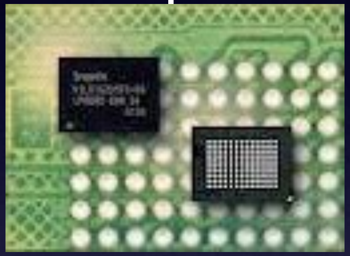

device with small memory

Attractive model for quantum computation

- initial quantum computers are likely to have few qubits
- captures fast processing of input, may cope with low coherence time
- goes beyond finite quantum automata

Streaming quantum algorithms

Advantage over classical

- Quantum finite automata: streaming algorithms with constant memory and time per symbol. Some are exponentially smaller than classical FA.

- Use exponentially smaller amount of memory for certain problems [LeG'06, GKKRdW'06]

Advantage for natural problems ?

- For Dyck(2), checking if an expression in two types of parentheses is well-formed ?

Quantum streaming complexity of Dyck(2) ?

Theorem [Jain, N. '11]

If a quantum protocol computes AI_n with probability $1 - \varepsilon$ on the uniform distribution, either

Alice reveals $\Omega(n / t)$ information about x, or

Bob reveals $\Omega(1 / t)$ information about k,

under the uniform distribution over 0-inputs, where t is the number of rounds.

- Specialized notion of information cost
- Connection to streaming algorithms breaks down
- Connection to *communication* complexity unclear
- Other notions: fixed above problems, but couldn't analyze

Results

$x = x_1 x_2 \ldots x_n$ Is $x_k = b$? k, $x[1, k\text{-}1]$, b

Theorem [N., Touchette '16]

* If a quantum protocol computes AI_n with probability $1 - \varepsilon$ on the uniform distribution, either

Alice reveals $\Omega(n / t^2)$ information about x, or

Bob reveals $\Omega(1 / t^2)$ information about k,

under the uniform distribution over 0-inputs, where t is the number of rounds.

* Any T-pass unidirectional quantum streaming algorithm for Dyck(2) uses $n^{1/2} / T^3$ qubits on instances of length n

Quantum information trade-off

- Uses a new notion, Quantum Information Cost [Touchette '15]
- High-level intuition and structure of proof similar to [Jain, N. '11], but new execution, uses new tools
- Overcomes earlier difficulties in analysis:
    - inputs to Alice and Bob are correlated
    - need to work with superpositions over inputs
    - superpositions leak information in counter-intuitive ways
- Develop a "fully-quantum" analogue of the "Average Encoding Theorem" [KNTZ'07, JRS'03]
- Use of tools needs special care

Lower bound for quantum streaming algorithms

- Define general model for quantum streaming algorithms: allows for measurements / discarding qubits (non-unitary evolution)
- Quantum Information Cost allows us to lift the [MMN'10] connection between streaming and low-information protocols, even for this general model
- Proof of information cost trade-off requires protocols with pure (unmeasured) quantum states
- QIC does not increase, when we transform protocols with intermediate measurements to those without

Main Result

$x = x_1 \, x_2 \ldots x_n$ Is $x_k = b$? $k,\ x[1, k\text{-}1],\ b$

Theorem [N., Touchette '16]

If a quantum protocol computes AI_n with probability $1 - \varepsilon$ on the uniform distribution, either

Alice reveals $\Omega(n / t^2)$ information about x, or

Bob reveals $\Omega(1 / t^2)$ information about k,

under the uniform distribution over 0-inputs, where t is the number of rounds.

Intuition behind proof

(2 classical messages, [JN'10])

$x = x_1\, x_2 \ldots x_n$

↓ output

$\longrightarrow M_A$

$\longleftarrow M_B$

$k,\ x[1, k\text{-}1],\ b$

Consider uniformly random X, K, let $B = X_K$ (0-input)

- Consider K in [n/2]. If M_A has $o(n)$ information about X, then it is nearly independent of X_L, $L > n/2$. Flipping Alice's L-th bit does not perturb M_A much.

- If M_B has o(1) information about K, then M_B is nearly the same, on average, for pairs $J \leq n/2$, $L > n/2$. Switching Bob's index from J to L does not perturb M_B much.

Consequences of Average Encoding Theorem [KNTZ'07, JRS'03]

Intuition continued...

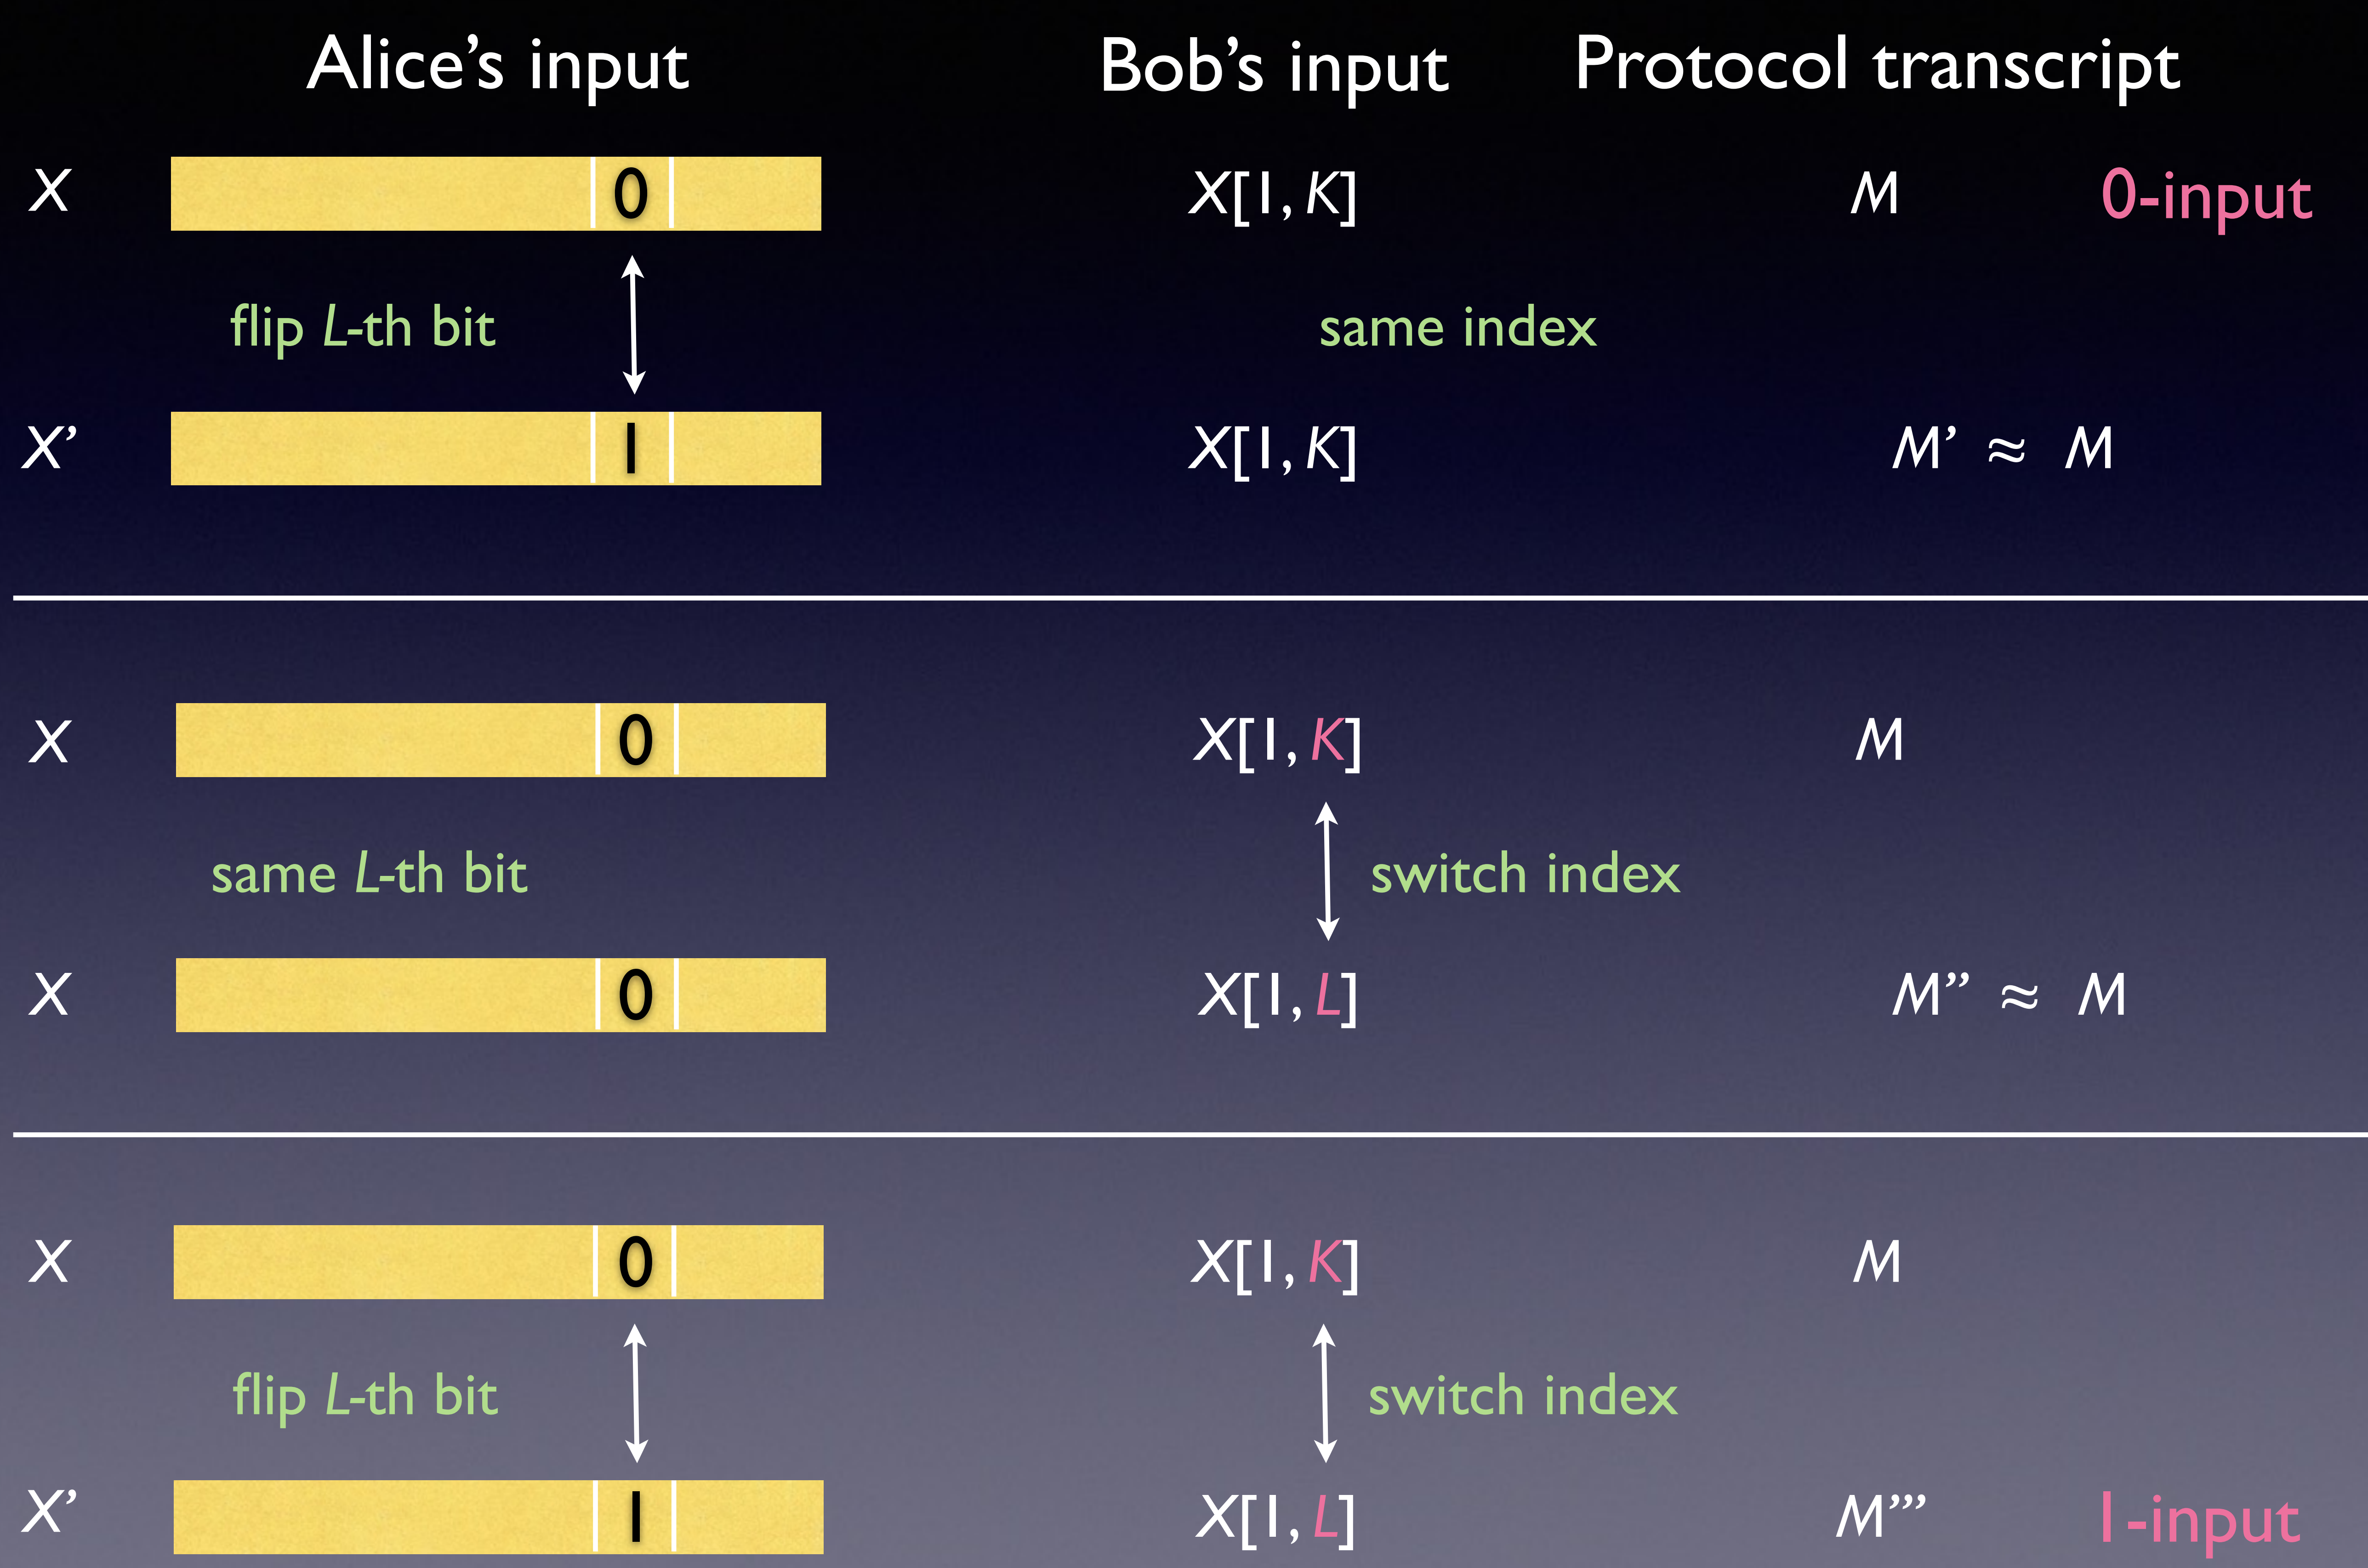

Finally...

	Alice's input		Bob's input		Protocol transcript	
X	0		$X[1, K]$		M	0-input
	flip L-th bit ↕		↕ switch index			
X'	1		$X[1, L]$		M'''	1-input

We have $M \approx M'$ and $M \approx M''$. Therefore, $M' \approx M''$ (triangle inequality)

Cut and paste lemma [BJKS'04]

In any (private coin) randomized protocol, the Hellinger distance between message transcripts on inputs (u,v) and (u',v') is the same as that between (u',v) and (u,v')

Therefore, $M \approx M'''$ and the (low-information) protocol errs.

Quantum case

(2 messages, both superpositions)

$x = x_1\, x_2 \ldots x_n$

$U_X \mid 0 \rangle$ →

← $\mid \psi \rangle = V_K\, U_X \mid 0 \rangle$

$k,\ x[1, k\text{-}1],\ b$

↓ output

Uniformly random X, K, let $B = X_K$ (0-input)

- Assume no party retains private qubits
- K in [n/2], $L > n/2$
- first message has o(n) information about X (given prefix), second message has little information about K (given X)

In this case, can use (quantum) mutual information, and Average Encoding Theorem [KNTZ'07, JRS'03]

Quantum case continued...

	Alice's input	Bob's input	Final protocol state
X	0	$X[1, K]$	$\vert \psi \rangle$ 0-input
	flip L-th bit ↕	same index	
X'	1	$X[1, K]$	$\vert \psi' \rangle \approx \vert \psi \rangle$

	Alice's input	Bob's input	Final protocol state
X	0	$X[1, K]$	$\vert \psi \rangle$
	same L-th bit	↕ switch index	
X	0	$X[1, L]$	$\vert \psi'' \rangle \approx \vert \psi \rangle$

	Alice's input	Bob's input	Final protocol state
X	0	$X[1, K]$	$\vert \psi \rangle$
	flip L-th bit ↕	↕ switch index	
X'	1	$X[1, L]$	$\vert \phi \rangle$ 1-input

Finally...

	Alice's input		Bob's input	Protocol state
X	0		$X[1, K]$	$\vert \psi \rangle$
	flip L-th bit ↕		↕ switch index	
X'	1		$X[1, L]$	$\vert \phi \rangle \approx \vert \psi \rangle$?

$$| \psi \rangle = V_K U_X | 0 \rangle , \quad | \psi' \rangle = V_K U_{X'} | 0 \rangle , \quad | \psi'' \rangle = V_L U_X | 0 \rangle$$

$$| \phi \rangle = V_L U_{X'} | 0 \rangle$$

$$| \varphi - \psi | \leq | \psi - \psi'' | + | \varphi - \psi'' |$$

$$\leq \delta + | V_L U_{X'} | 0 \rangle - V_L U_X | 0 \rangle |$$

$$= \delta + | V_K U_{X'} | 0 \rangle - V_K U_X | 0 \rangle |$$

$$= \delta + | \psi - \psi' | \leq 2 \delta$$

Details omitted

- Alice and Bob may maintain private workspace, communicate over more rounds
- Need to use QIC to quantify information, work with superpositions over inputs
- Use "superposed average encoding theorem", building on a 2015 breakthrough by Fawzi-Renner
- Perturbation of message due to switching of input depends on the number of rounds
- Hybrid argument conducted round by round à la [JRS'03]
- Leads to round-dependant trade-off
- Trade-off can be strengthened using ideas from [Lauriere and Touchette'16], can then work with Average Encoding Theorem

Final remarks

- Established a trade-off for quantum information cost for Augmented Index
- Round dependence probably an artefact of the proof; eliminating this is related to question about Disjointness
- Implies a space lower bound for streaming algorithms for Dyck(2): matches classical case, up to round-dependence
- Tools may be useful more generally in quantum communication complexity