What Can be Observed Locally?

Round-based Models of Quantum Distributed Computing

Cyril Gavoille
Adrian Kosowski
LaBRI - University of Bordeaux

Marcin Markiewicz
Institute of Theoretical Physics and Astrophysics
University of Gdańsk

Outline

- What does ***quantum*** mean?
  - Some intuition
  - Some definitions
- How does Quantum Information help?
  - Quantum Computing – centralised models
  - Quantum Communication Complexity
- Where does ***locality*** come into play?
  - Locality in Computer Science *vs*. Locality in Physics
  - Quantum extensions of Linial's *LOCAL* model
  - Proving lower bounds on the round complexity of problems

GOAL

A quantization of the *LOCAL* model:

Creating a quantum model which:

- **when restricted to classical states is precisely the *LOCAL* model**
- **captures the same principles of locality.**

A simple experiment in a classical world…

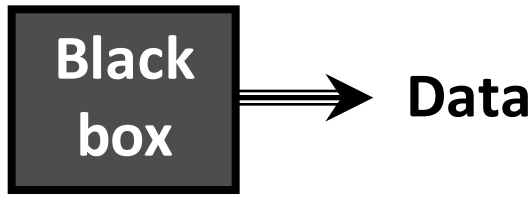

1. The observables

- What do we assume about the structure of the data?
  Let's say we know that the data is a pair of bits $(b_1,b_0) \in \Upsilon = \{00, 01, 10, 11\}$.
- What characteristics of the data are measurable?
  Let's say we can measure the value of bits b_1 and b_0 directly.

2. Measurement of the state

- What does the (randomized) algorithm in the black box produce?
  Identification: run an experiment many times independently, measuring $A = 2b_1 + b_0$ each time, obtain a probability distribution of values…

Describing a physical system

- ***Let's say the box flips a coin and outputs 01 or 10***. We have the state μ:
  $\mu(\underline{00}) = \mu(\underline{11}) = 0$, $\mu(\underline{01}) = \mu(\underline{10}) = ½$
- Observables: random variables <=> a commutative matrix algebra over complex numbers

$$\begin{array}{c} \\ \mathbf{B}_0 = \end{array}\begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & & & \\ & 1 & & \\ & & 0 & \\ & & & 1 \end{bmatrix} \end{array} \quad \begin{array}{c} \\ \mathbf{B}_1 = \end{array}\begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & & & \\ & 0 & & \\ & & 1 & \\ & & & 1 \end{bmatrix} \end{array} \quad \begin{array}{c} \\ \mathbf{A} = 2\mathbf{B}_1 + \mathbf{B}_0 = \end{array}\begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & & & \\ & 1 & & \\ & & 2 & \\ & & & 3 \end{bmatrix} \end{array}$$

- Permissable probabilistic measures are described by linear functionals over the defined algebra of observables

$\mathbf{E}_\mu A = \sum_{x\in\Xi} [A(x)*\mu(x)]$ <=> $\mathbf{E}_\mu \mathbf{A} = \mathrm{Tr}\,(\mathbf{A}\,\boldsymbol{\mu})$

(the trace is the sum of elements on the diagonal)

$$\begin{array}{c} \\ \boldsymbol{\mu} = \end{array}\begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & & & \\ & 1/2 & & \\ & & 1/2 & \\ & & & 0 \end{bmatrix} \end{array}$$

Describing a physical system

What is a measurement?

- Recall that we were measuring an observable **A** in state $\boldsymbol{\mu}$

$$\mathbf{A} = 2\mathbf{B}_1 + \mathbf{B}_0 = \begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & & & \\ & 1 & & \\ & & 2 & \\ & & & 3 \end{bmatrix} \end{array} \qquad \boldsymbol{\mu} = \begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & & & \\ & 1/2 & & \\ & & 1/2 & \\ & & & 0 \end{bmatrix} \end{array}$$

- The expected result of the measurement was given as: Tr $(\mathbf{A}\,\boldsymbol{\mu}) = 1.5$
- The possible outcomes are {0, 1, 2, 3} with probabilities {0, 1/2, 1/2, 0}, resp.
  - The outcomes are the eigenvalues λ_i of the matrix **A**… ($\mathbf{A} = \sum \lambda_i \mathbf{P}_i$)
  - The probability of obtaining outcome λ_i is exactly Tr $(\mathbf{P}_i\,\boldsymbol{\mu})$
- What changes in the quantum case?
  We allow **A** to be any complex-valued matrix with positive (real) eigenvalues.

Introduction

The Quantum Framework

- As computer scientists, we will find the following intuition useful:

The quantum framework is a generalization of classical probability

  - quantum algorithms are more powerful than randomized algorithms
  - quantum information can be manipulated in ways in which classical information cannot

Why extension of probability is required?

The problem with our universe…

- It is possible to perform a physical experiment in which we look at 4 characteristics of a simple system, and obtain marginal distributions for which there does not exist a joint distribution, in ***any*** probabilistic space.
  - So called ”violation of Bell’s Theorem”, first verified by Aspect (1982).
- Quantum Mechanics has to rely on an extension of the classical framework

Describing a physical system

What properties must a quantum state fulfill?

- Must be a density matrix (positive spectrum, trace normalised to 1)
- Two examples of valid states (density matrices):

$$\begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \boldsymbol{\mu}_1 = \begin{bmatrix} 0 & & & \\ & 1/2 & & \\ & & 1/2 & \\ & & & 0 \end{bmatrix} \end{array} \qquad \begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \boldsymbol{\mu}_2 = \begin{bmatrix} 0 & & & \\ & 1/2 & 1/2 & \\ & 1/2 & 1/2 & \\ & & & 0 \end{bmatrix} \end{array}$$

- Do $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ describe the same state? [Recall that: $\mathbf{E}_{\mu}\mathbf{A} = \mathrm{Tr}\,(\mathbf{A}\,\boldsymbol{\mu})$]
  Depends on what characteristics of the system are observable…
  - For the classical example with diagonal observables only – same state
  - For a richer class of quantum observables – these are distinct states…
    - The state $\boldsymbol{\mu}_2$ has no good classical interpretation!

Describing a physical system

Dirac's bra-ket notation for pure states

- A state $\boldsymbol{\mu}$ is called projective if $\boldsymbol{\mu} = \psi^{+}\psi$ for some row vector ψ
  - the cross ($^{+}$) denotes Hermitian transpose – transpose & conjugate
  - projective states are equivalent to so-called pure states in this context

$$\boldsymbol{\mu}_2 = \begin{array}{c} \begin{array}{cccc} 00 & 01 & 10 & 11 \end{array} \\ \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{array} = \begin{bmatrix} 0 \\ 1/\sqrt{2} \\ 1/\sqrt{2} \\ 0 \end{bmatrix} \begin{array}{cccc} 00 & 01 & 10 & 11 \\ [0 & 1/\sqrt{2} & 1/\sqrt{2} & 0] \end{array}$$

- It is often more convenient to work on such vectors ψ, especially when using tensor products. A basis vector is usually written as a |ket>:

$$\psi = \begin{array}{cccc} 00 & 01 & 10 & 11 \\ [0 & 1/\sqrt{2} & 1/\sqrt{2} & 0] \end{array} = 1/\sqrt{2}\ |01\rangle + 1/\sqrt{2}\ |10\rangle$$

What is a quantum bit?

- A classical bit: 0 or 1
- A probabilistic classical bit: $(p_0 \ p_1)$
- A quantum bit (or qubit):

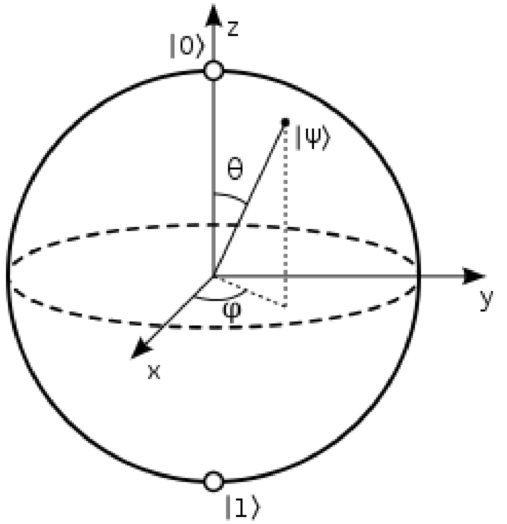

$$\alpha \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} + \beta \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} = \alpha|\mathbf{0}> + \beta|\mathbf{1}>$$

... where α,β are complex numbers

- State of a quantum sytem ψ is denoted by |ψ> (bra/ket notation)

  (a 1-qubit or a *n*-qubit register)

Quantum Operators

- Operators on *n*-qubit system are represented by 2^nx2^n complex-valued matrices
- There are restrictions on the possible operators, usually: **unitary matrix**.
- In particular, there is no operator M such that M.|Ψ0> = |ΨΨ>

  => No Cloning Property

- Classical operator (NOT, AND ...) can be converted into algebraic operators (by adding extra wires)
- Like classical *n*-bit operations, quantum operators can be decomposed as combinations (products and tensor products) of 1-qubit operators (gates)

  => Quantum Universal Turing Machine

A centralised quantum computer

Quantum circuits: the set-up

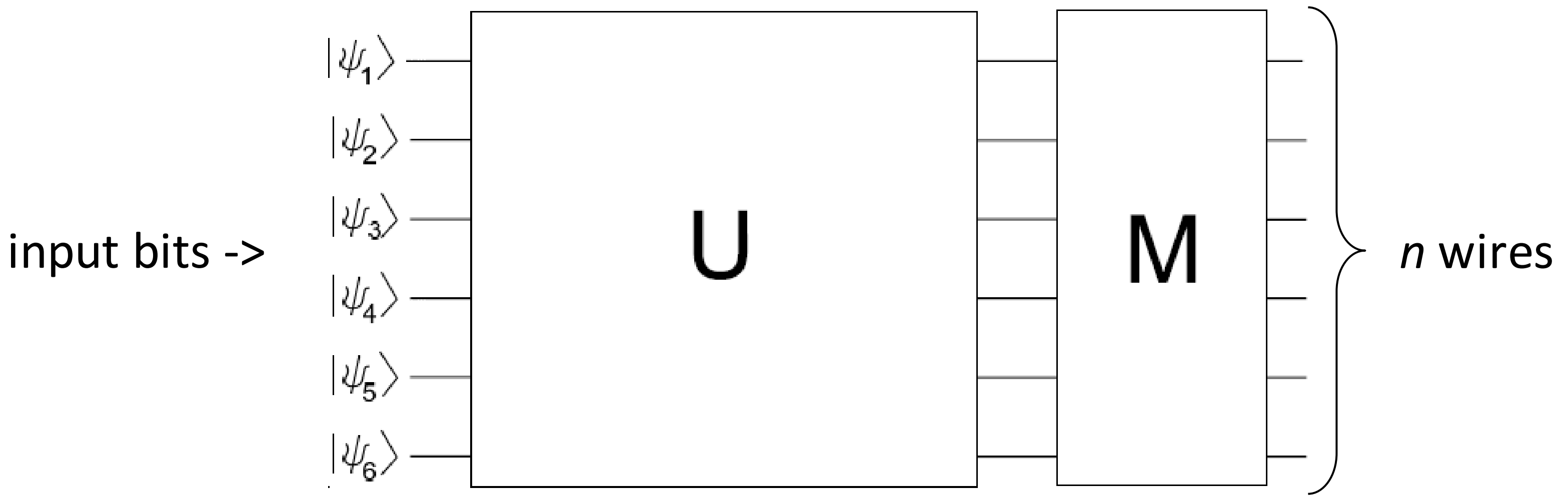

- Goal: transform a *k*-bit input vector into an *n*-bit output vector
  - Encoding classical input the quantum way
  - **U** – transforming the quantum information (quantum operations)
  - **M** – performing a measurement to obtain classical output

A centralised quantum computer

Transforming the data: what is feasible?

- Any unitary matrix can be built up from single-input gates, and the two-input controlled-not (CNOT) gate

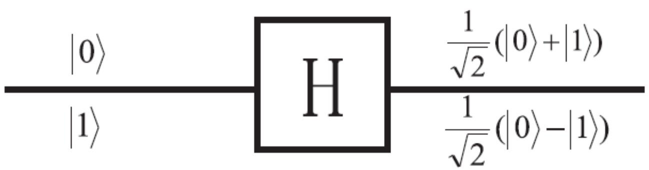

(a) The Hadamard gate

$$H_1 = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

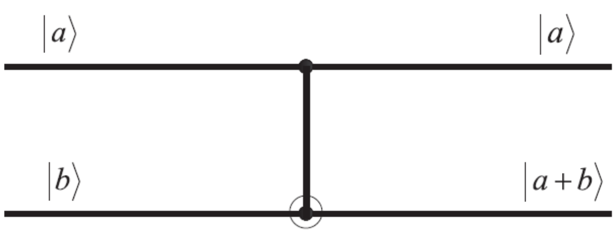

(b) Quantum controlled-NOT

$$\mathrm{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

- Certain operations cannot be performed, e.g. ”qubit copying”.

A centralised quantum computer

Transforming the data: what's the complexity of operation U?

- Building up the system from elementary bricks.

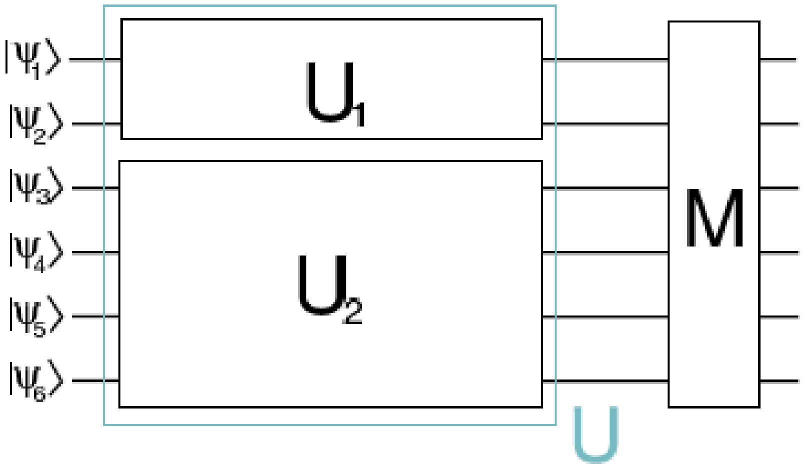

$\mathbf{U} = \mathbf{U}_1 \otimes \mathbf{U}_2$

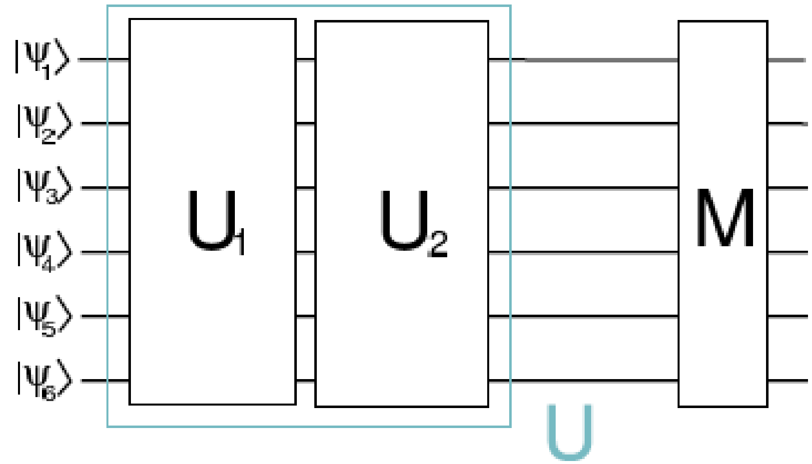

$\mathbf{U} = \mathbf{U}_2\mathbf{U}_1$

- For both types of combinations, the complexity measure is subadditive:

$$C(\mathbf{U}) \leq C(\mathbf{U}_1) + C(\mathbf{U}_2)$$

- Elementary gates acting on spaces of size O(1) are assumed to have complexity O(1).

Quantum Distributed Computing

Why should quantum information help?

- **Negative evidence**: even when Alice and Bob share entanglement, they cannot do any magic

  (E.g. no way to exchange information without sending messages)
- **Negative evidence**: **Holevo's** theorem – the usable information content (entropy) of an *n*-qubit state is not greater than that of an *n*-bit string.
  - so, does it make sense to send qubits at all?
- But: it turns out that
  Quantum information sometimes reduces communication complexity
  - One possible explanation: Information is no longer encoded at specific locations. ***The state is a global property of the system*.**
- First examples:
  - Grover's O(√*n*)-time search algorithm, 1997
  - Cleve & Buhrman, 1997 – simple 3-party proof-of-concept example

Example (Cleve & Buhrman, 1997)

Problem definition

- Three parties, Alice, Bob, and Carol, are given an *n*-bit input string, each (strings *a*, *b*, *c*, respectively).
- It is known that for all *n* indices, the bits fulfill the condition: $a_i \oplus b_i \oplus c_i = 1$
- Goal: Alice is to compute the value of $a_1b_1c_1 \oplus a_2b_2c_2 \oplus \ldots \oplus a_nb_nc_n$

Theorem. *Any classical protocol requires communication of at least 3 bits.*

Quantum solution

- We do not change the communication capabilities of the system – classical messages (classical bits) only.
- We allow Alice, Bob and Carol to preshare an ***entangled*** state:
  1/2 (|001> + |010 > + |100 > − |111 >) // repeated *n* times
- Now, the problem can be solved using 2 communicated bits in total (Bob sends Alice 1 bit, Carol sends Alice 1 bit.)

Example (Cleve & Buhrman, 1997)

Details

- Each party *p* transforms its *i*-th qubit (q_i) depending on the values of the *i*-th input bit (x_i).

$$
\begin{aligned}
&\textbf{for each } i \in \{1, \ldots, n\} \textbf{ do}\\
&\qquad \textbf{if } x_i^p = 0 \textbf{ then apply } H \textbf{ to } \mathrm{q}_i^p\\
&\qquad \textbf{measure } \mathrm{q}_i^p \textbf{ yielding bit } s_i^p\\
&s^p \leftarrow s_1^p + \cdots + s_n^p
\end{aligned}
$$

$$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

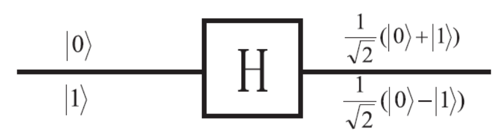

- Each party other than Alice transmits its bit *s* to Alice.
- Alice returns $s^A \oplus s^B \oplus s^C$ as output.

The Question of Locality

The classical *LOCAL* model

Assumptions of the *LOCAL* model

- The distributed system consists of a set of processors V, $|V|=n$
- The system operates in synchronous rounds
- No faults are present
- The system input is encoded as a ***labeled*** graph:
  - edge set E; $G=(V,E)$
  - node labels $x(v)$, for $v \in V$
- The result of computations is given through local variables $y(v)$, for $v \in V$
- *Messages exchanged in each round may have unbounded size*
- *The computational capabilities of each node are unbounded*
- As a rule, we will assume that nodes have unique identifiers

Extending the *LOCAL* model

Quantum extensions

- **System initialization** (before the input is set)
  - by default: all the processors have an identical starting state
  - +S: the algorithm may predefine any global separable (=classical) state as a starting state of the system
  - +E: the algorithm may predefine any global entangled (=quantum) state as a starting state of the system

- **Communication capabilities**
  - by default: the processors communicate by exchanging classical messages (bits)
  - +Q: in each round, the processors can communicate by exchanging quantum information (qubits)

The *LOCAL+E* model

How much does the +E extension help?

- **+E: Entangled initial state**
  - allows us to take full advantage of quantum capabilities of the system
- Proof-of-concept "Mod 4" problem showing that +E does help:
  Variant of famous Greenberger-Horne-Zeilinger (GHZ) experiment
  - V consists of 3 nodes $\{v_1, v_2, v_3\}$, whereas E is empty
  - Each node has an input label $x_i \in \{0,1\}$ provided $(x_1 + x_2 + x_3) \in \{0,2\}$
  - **Goal**: output labels $y_i \in \{0,1\}$ must be such that:

$$2(y_1 + y_2 + y_3) \equiv (x_1 + x_2 + x_3) \bmod 4$$

  - cannot be solved with Pr > ¾ in (classical) LOCAL+S model, in any time
  - can be solved deterministically in 0 rounds with the +E extension (pre-shared GHZ state |000>+|111>)

Mod 4 problem – a conceptual look

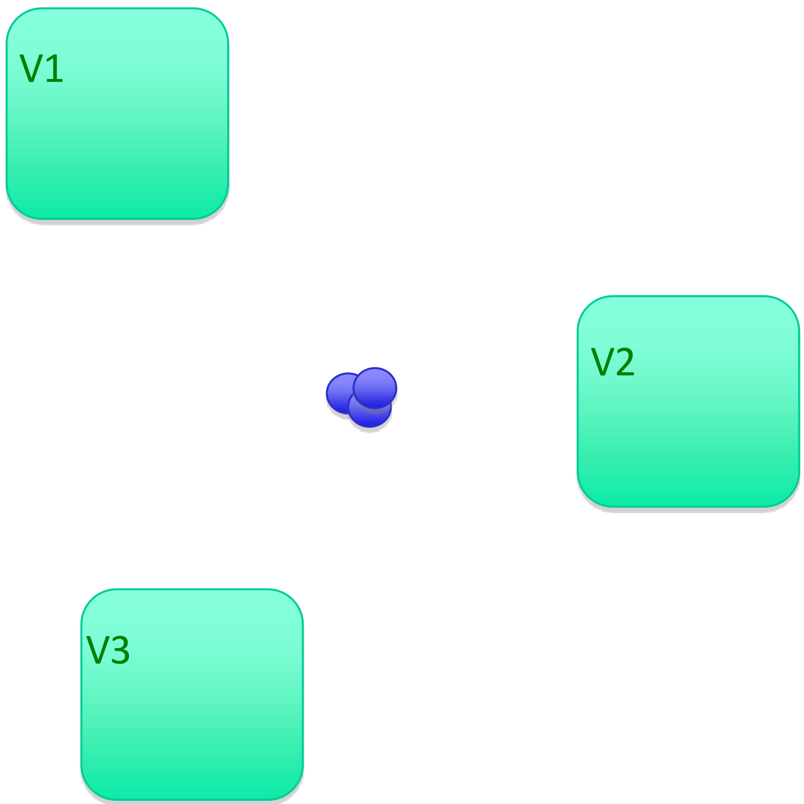

Mod 4 problem – a conceptual look

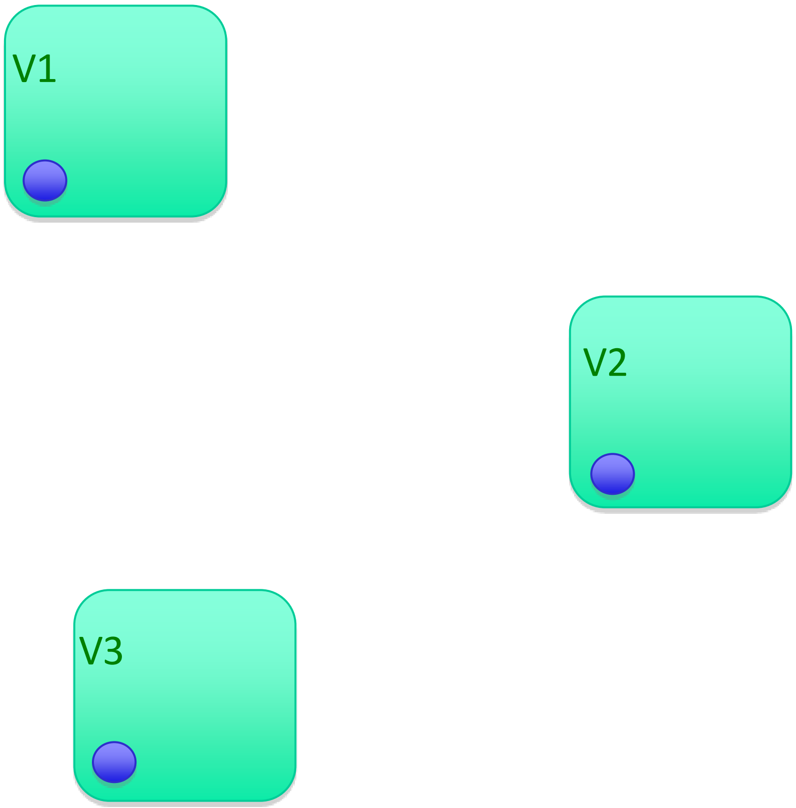

Mod 4 problem – a conceptual look

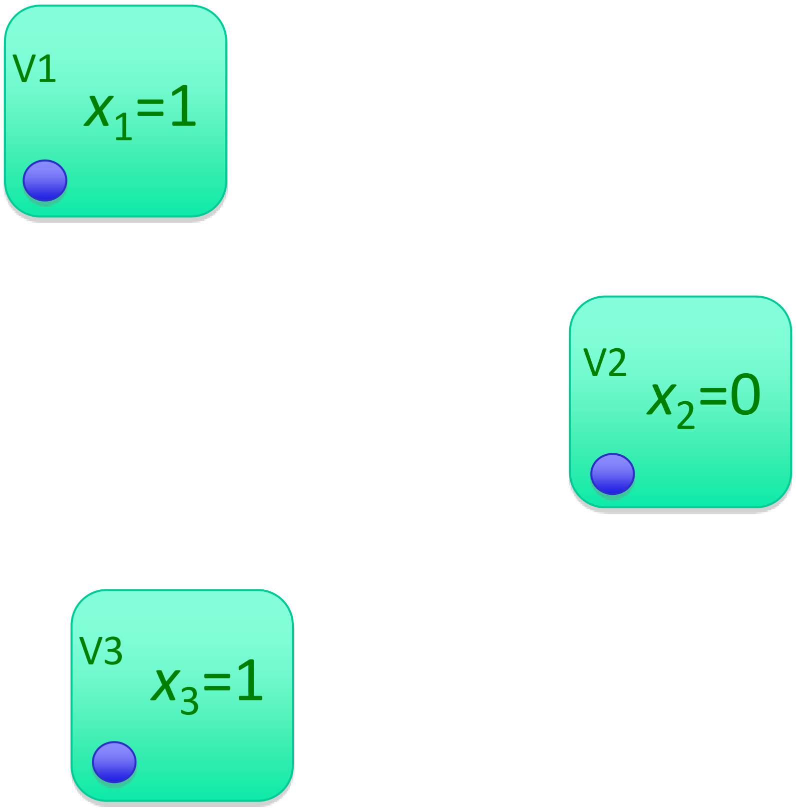

Mod 4 problem – a conceptual look

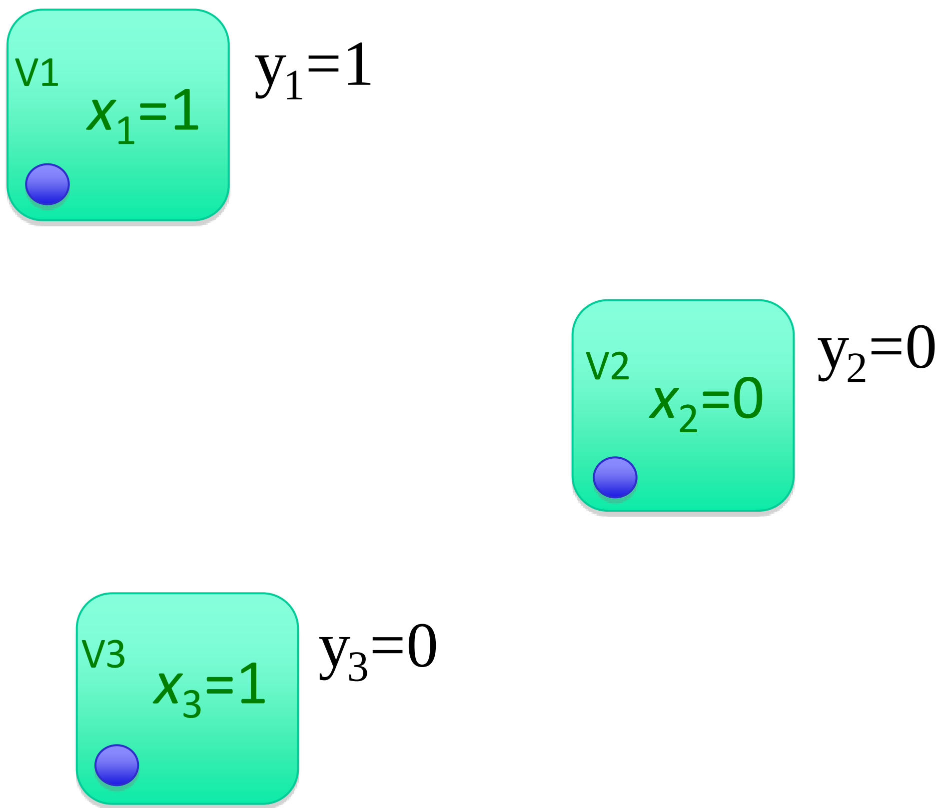

Mod 4 problem – a conceptual look

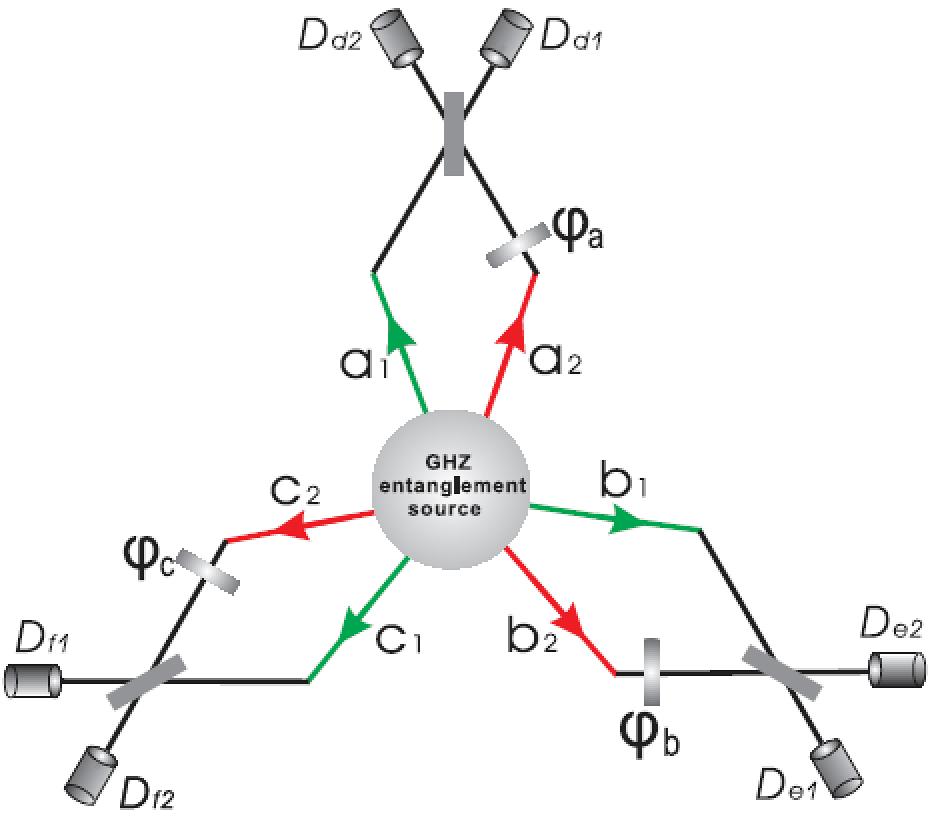

The ΛΟΧΑΛ+E model

Outcome of a quantum algorithm for the "Mod 4" problem

Input (x_1, x_2, x_3)	Probability p^i	Output (y_1^i, y_2^i, y_3^i)
(0, 0, 0)	1/4	(0, 0, 0)
	1/4	(0, 1, 1)
	1/4	(1, 0, 1)
	1/4	(1, 1, 0)

Input (x_1, x_2, x_3)	Probability p^i	Output (y_1^i, y_2^i, y_3^i)
(0, 1, 1) or (1, 0, 1) or (1, 1, 0)	1/4	(1, 1, 1)
	1/4	(1, 0, 0)
	1/4	(0, 1, 0)
	1/4	(0, 0, 1)

$$2(y_1 + y_2 + y_3) \equiv (x_1 + x_2 + x_3) \bmod 4$$

A comparison of the computational power of quantum models

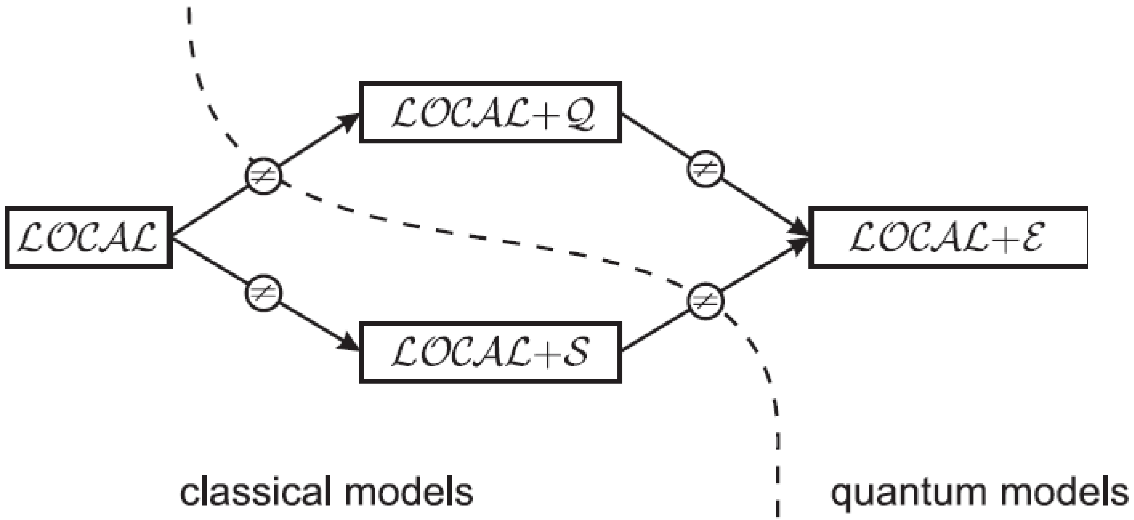

The quantum models are more powerful than the classical ones.

Do they have any natural limits (lower time bounds)?

A comparison of the computational power of quantum models

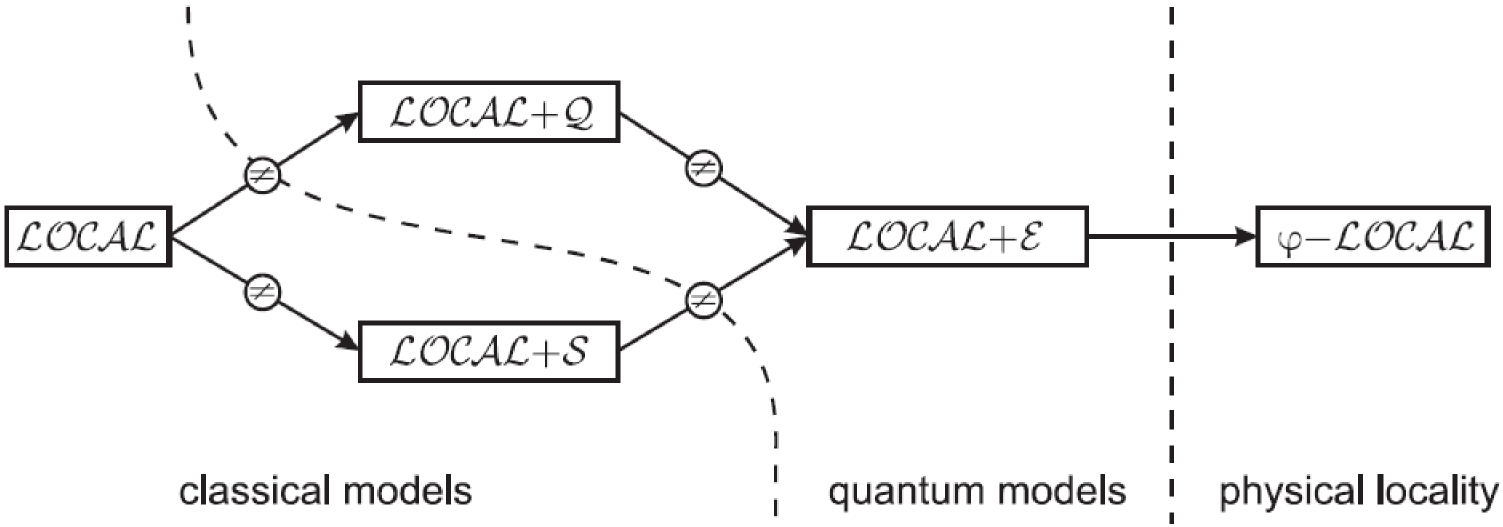

The quantum models are more powerful than the classical ones.

Do they have any natural limits (lower time bounds)?

The meaning of locality

Understanding of locality in the *LOCAL* model

- **each node builds up its view during the execution of the algorithm**
  - after t rounds, $VIEW_t(Gx, v)$ describes the distance-t neighbourhood of v in the labeled input graph Gx
- when considering deterministic algorithms, an output vector $\boldsymbol{y}$ can be reached in t rounds if and only if there exists a function f such that:

$$y(v) = f\,(\,VIEW_t(Gx, v)\,), \qquad \text{for all } v \in V$$

- this intuition can be extended to allow for randomized algorithms.
- no similar complete characterization is known for quantum approaches
  - doing it precisely would give a nice result on the capabilities of quantum operations (completely positive maps)
- **However**: we know of a weaker, but still view-based, bound on the computational power of any quantum algorithm.

The meaning of locality

Physical locality: the φ-*LOCAL* model

- **Thesis.** Locality is violated if and only if, based on the available output data, we can conclusively verify that after t rounds:
  some subset S of processors was affected by input data initially localized outside its view, which is $VIEW_t(Gx, S) := \bigcup_{v \in S} VIEW_t(Gx, v)$.
- The preservation of locality should be interpreted in a probabilistic way:
  - consider the outcome of an algorithm after t rounds; for the subset S, we look at the probability p of obtaining any given output vector $\mathbf{y}[S]$
  - if two inputs differ only by edges/label located outside $VIEW_t(Gx^{(i)}, S)$, then this probability p must necessarily be the same for both inputs
  - (otherwise, we would be able to detect this remote difference in the input by performing many parallel executions of our algorithm)
- φ-*LOCAL* is provably not less powerful than the quantum models

The meaning of locality

Example: why is the "Mod 4" problem in φ-*LOCAL*?

Input (x_1, x_2, x_3)	Probability p^i	Output (y_1^i, y_2^i, y_3^i)
(0, 0, 0)	1/4	(0, 0, 0)
	1/4	(0, 1, 1)
	1/4	(1, 0, 1)
	1/4	(1, 1, 0)

Input (x_1, x_2, x_3)	Probability p^i	Output (y_1^i, y_2^i, y_3^i)
(0, 1, 1) or (1, 0, 1) or (1, 1, 0)	1/4	(1, 1, 1)
	1/4	(1, 0, 0)
	1/4	(0, 1, 0)
	1/4	(0, 0, 1)

- we consider the above solution (obtained by the quantum algorithm), and one by one all possible sets *S.*
  - for example, let $S=\{v_1\}$; since the graph is empty, $VIEW_t(Gx, S) = \{v_1\}$
  - what are the probabilities of particular outputs?
    - in this case, regardless of *Gx*: $\Pr[y_1=0] = ½$ and $\Pr[y_1=1] = ½$
  - so, these probabilities are not affected by the values of x_2, x_3, and the φ−*LOCAL* condition is not violated.

Lower time bounds in quantum models

So, what lower bounds can be proved in φ-*LOCAL*?

- Most proofs of lower time bounds which rely on view-based arguments will hold in φ-*LOCAL* (and hence also all the quantum models)
  - The problem of finding a maximal independent set in the system graph requires $\Omega(\sqrt{(\log n / \log\log n)})$ rounds to solve [Kuhn, Moscibroda, Wattenhofer, 2004]
  - The problem of finding a locally minimal (greedy) coloring of the system graph requires $\Omega(\log n / \log\log n)$ rounds to solve [G., Klasing, K., Navarra, Kuszner, 2009]
  - The problem of finding a spanner with $O(n^{1+1/k})$ edges requires $\Omega(k)$ rounds to solve [Elkin 2007; Derbel et al. 2008]
- What about Linial's famous $\Omega(\log^* n)$ bound on $(\Delta+1)$-coloring?
  - The neighbourhood-graph technique does not work in φ-Λ*OCAL* …

Lower time bounds in quantum models

Example: time required to 2-color the even ring

- In the *LOCAL* model, $n/2 - 1$ rounds are required and sufficient
  - simpler version of the same neighbourhood graph technique
- In φ-*LOCAL*, $\lceil n-2 \rceil / 4$ rounds are required and sufficient
- Sketch of lower bound
  - let $t < \lceil n-2 \rceil / 4$, there will be at least two nodes *u* and *v* of the ring whose views are still disjoint
  - let $S = \{u,v\}$;
  - the color values of *u* and *v* are necessarily the same if these vertices are at an even distance, and odd otherwise
    - there exist corresponding input graphs $Gx^{(1)}$ and $Gx^{(2)}$ with odd and even distance between *u* and *v*, respectively
  - but the difference cannot be detected based on the local views of *u* and *v.*

Lower time bounds in quantum models

Is it possible to design a real quantum routine for 2-coloring C_6 in 1 round?

in the *LOCAL* model, 2 rounds are required and sufficient

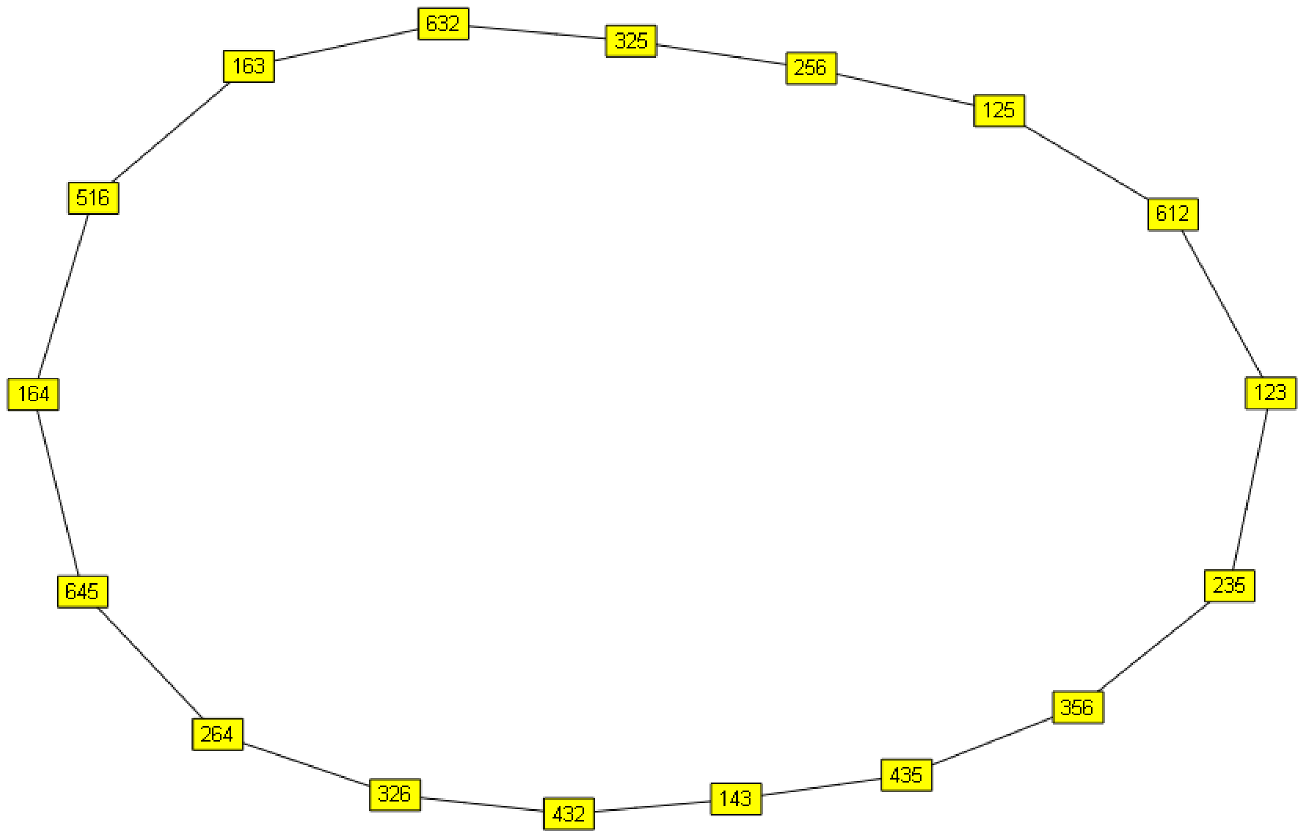

in the φ-*LOCAL* model, 1 round is required and sufficient

Lower time bounds in quantum models

Some open problems:

- Can quantum distributed algorithms be designed for any combinatorial problems of significance to practice or theory?
- How many rounds are required to 3-color the ring in the studied quantum models and in φ-*LOCAL*?
- What is the lower time bound on the $(\Delta+1)$-coloring problem in quantum models? (currently all we know is that we need at least one round…)
- Is it possible to design a real quantum routine for 2-coloring C_6 in 1 round? (in the φ-*LOCAL* model 1 round is required and sufficient)
- Does *LOCAL*+E = φ-*LOCAL*?

Thank You!