

Beyond Classical Local Models in Distributed Computing

Cyril Gavoille

LaBRI, University of Bordeaux

44th Ecole de Printemps en Informatique Théorique (EPIT)

- Distributed Computing -

Île de Porquerolles, May 15-19, 2017

Beyond The LOCAL Model

Goal. Study the locality of a problem when nodes have access to exotic resources, like entangled qubits or others.

Reminder:

- Nodes have IDs
- Run in synchronous rounds
- No message size limits
- No local computational power limits

Beyond The LOCAL Model

Goal. Study the locality of a problem when nodes have access to exotic resources, like entangled qubits or others.

Reminder:

- Nodes have IDs
- Run in synchronous rounds
- No message size limits
- No local computational power limits

Fact. 3-coloring an n -node ring requires $\frac{1}{2} \log^* n$ rounds.
(Similar for a path and for MIS computation.)

Extending the LOCAL Model

Extensions:

- 1 Shared-randomness
- 2 Quantum resources
- 3 Arbitrary resources preserving physical locality

Extending the LOCAL Model

Extensions:

- 1 Shared-randomness
- 2 Quantum resources
- 3 Arbitrary resources preserving physical locality

Fact. It is not known if 3-coloring an n -node ring can be done faster than $\frac{1}{2} \log^* n$ rounds if nodes have quantum resources

Extending the LOCAL Model

Extensions:

- 1 Shared-randomness
- 2 Quantum resources
- 3 Arbitrary resources preserving physical locality


Fact. It is not known if 3-coloring an n -node ring can be done faster than $\frac{1}{2} \log^* n$ rounds if nodes have quantum resources

There are proof-of-concept problems showing that it may help.

Shared-Randomness

Reminder. Randomness does not help for 3-coloring a ring.

Task. Select in 0 rounds two adjacent nodes in $K_n \setminus \{e\}$. Exactly two adjacent nodes must output **select**, the others **unselect**.



Shared-Randomness

Task. Select in 0 rounds two adjacent nodes in $K_n \setminus \{e\}$. Exactly two adjacent nodes must output **select**, the others **unselect**.

Cannot be solved in the classical LOCAL model with $\Pr > 1/e$.

[One can tends to $1/e$ by always selecting node with ID=1 and each other with $\Pr = (n - 2)/(n - 1)$, if $n \geq 6$]

Shared-Randomness

Task. Select in 0 rounds two adjacent nodes in $K_n \setminus \{e\}$. Exactly two adjacent nodes must output **select**, the others **unselect**.

Cannot be solved in the classical LOCAL model with $\Pr > 1/e$.

[One can tends to $1/e$ by always selecting node with ID=1 and each other with $\Pr = (n - 2)/(n - 1)$, if $n \geq 6$]

With shared-randomness, each node can select any edge of K_n uniformly at random in a coordinated manner. This is correct with $\Pr = 1 - 1/\binom{n}{2} = 1 - O(1/n^2)$.

Quantum Extension

Nodes may share arbitrary entangled states (before inputs come) ... use quantum communication channels with teleportation effects and others.

Quantum Extension

Nodes may share arbitrary entangled states (before inputs come) ... use quantum communication channels with teleportation effects and others.

To understand why quantum extensions is strictly stronger than shared-randomness extensions, we need to introduce quantum information background ...

Quantum Physics

(for computer scientists)



Causality & Locality

Two fundamental principles in physics: **causality** and **locality**.

Causality & Locality

Two fundamental principles in physics: **causality** and **locality**.

Reproducible experiments of Alain Aspect in '82 show that one of the two principles (at least) does not hold. Bell's inequalities are violated ...

Causality & Locality

Two fundamental principles in physics: **causality** and **locality**.

Reproducible experiments of Alain Aspect in '82 show that one of the two principles (at least) does not hold. Bell's inequalities are violated ...

Quantum effects/non-locality: a quantum particle (the support of information) can be in several places at the same time ...

Causality & Locality

Two fundamental principles in physics: **causality** and **locality**.

Reproducible experiments of Alain Aspect in '82 show that one of the two principles (at least) does not hold. Bell's inequalities are violated ...

Quantum effects/non-locality: a quantum particle (the support of information) can be in several places at the same time ...

Unlike relativity theory, quantum physics do not logically follows from easy-to-validate axioms, like the invariance principle (laws are the same under any coordinates axis) and causality. BTW, invariance and causality imply an upper bound speed limit for light.

Unit of Information

Deterministic bit B : represented by a value $B \in \{0, 1\}$.

Unit of Information

Deterministic bit B : represented by a value $B \in \{0, 1\}$.

Probabilistic bit P : represented by a vector of two reals $\begin{pmatrix} p \\ q \end{pmatrix}$ where p is the probability to get 0 and q to get 1. Need $0 \leq p, q \leq 1$ and $p + q = 1$.

Denoting by 0 the vector $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ et by 1 the vector $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$, one can represent P as:

$$P = p \begin{pmatrix} 1 \\ 0 \end{pmatrix} + q \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} p \\ q \end{pmatrix}$$

Unit of Information


Quantum bit (or qubit) Q : represented by a vector of two complexes $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$ where $|\alpha|^2$ is the probability to get 0 and $|\beta|^2$ to get 1. Need $|\alpha|^2 + |\beta|^2 = 1$.

Denoting by 0 the vector $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ et by 1 the vector $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$, one can represent P as:


$$Q = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$$

Qubit


A qubit is implemented by a physical particle (photon, electron, atome). The **measure** of a qubit is a physical measurement of an **observable**, a property like: polarisation, spin, electrical charge, energy level, speed, impulsion, ...




Superposition & Interference



Superposition even for a single particle ...



Reality?



Entanglement of Qubit Systems

A system of two probabilist bits $\begin{pmatrix} p_1 \\ q_1 \end{pmatrix}$ and $\begin{pmatrix} p_2 \\ q_2 \end{pmatrix}$ represented by a tensor product.

$$\begin{aligned} \begin{pmatrix} p_1 \\ q_1 \end{pmatrix} \otimes \begin{pmatrix} p_2 \\ q_2 \end{pmatrix} &= \begin{pmatrix} p_1 \cdot \begin{pmatrix} p_2 \\ q_2 \end{pmatrix} \\ q_1 \cdot \begin{pmatrix} p_2 \\ q_2 \end{pmatrix} \end{pmatrix} = \begin{pmatrix} p_1 p_2 \\ p_1 q_2 \\ q_1 p_2 \\ q_1 q_2 \end{pmatrix} \\ &= p_1 p_2 \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + p_1 q_2 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} + q_1 p_2 \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} + q_1 q_2 \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} \end{aligned}$$

Entanglement of Qubit Systems

A system (A, B) composed of two qubits represented as linear combinations of four unit vectors:

$$(A, B) = \begin{pmatrix} c_{00} \\ c_{01} \\ c_{10} \\ c_{11} \end{pmatrix} = c_{00} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + c_{01} \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} + c_{10} \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} + c_{11} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

where $c_{ij} \in \mathbb{C}$ and $\sum_{ij} |c_{ij}|^2 = 1$, the probability to observe $A = i$ et $B = j$ being controlled by $|c_{ij}|^2$.

Entangled vs. Separable States

The state of a system (A, B) is **separable** if $(A, B) = A \otimes B$.
Otherwise it is **entangled**.

$$(A, B) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

Operators

The state of a system S of n qubits is denoted $|S\rangle$ which is a vector in \mathbb{H}^{2^n} . Operations on state vector is a **reversible** unit matrix of $\mathbb{H}^{2^n \times 2^n}$.

For $n = 1$,

$$\text{NOT} := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \text{NOT} \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

Operators

The state of a system S of n qubits is denoted $|S\rangle$ which is a vector in \mathbb{H}^{2^n} . Operations on state vector is a **reversible** unit matrix of $\mathbb{H}^{2^n \times 2^n}$.

For $n = 1$,

$$\text{NOT} := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \text{NOT} \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

$$R(\theta) := \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

Operators

The state of a system S of n qubits is denoted $|S\rangle$ which is a vector in \mathbb{H}^{2^n} . Operations on state vector is a **reversible** unit matrix of $\mathbb{H}^{2^n \times 2^n}$.

For $n = 1$,

$$\text{NOT} := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \text{NOT} \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

$$R(\theta) := \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}$$

$$\sqrt{\text{NOT}} := \frac{1}{2} \begin{pmatrix} 1+i & 1-i \\ 1-i & 1+i \end{pmatrix}, \quad \sqrt{\text{NOT}} \cdot \sqrt{\text{NOT}} = \text{NOT}$$


No Cloning & No Erasing

There is no operator U that for a qubit A that makes:



$$U \cdot |A, 0\rangle = |A, A\rangle \quad \text{or} \quad U \cdot |A\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

On Bell's Inequalities


Using a *Gedankenexperiment* invented by David Bohm (1951), John S. Bell (1964) showed that the nonexistence of properties before measurement is a direct consequence of the quantitative numerical predictions of the quantum theory. Alain Aspect (1982) showed that the quantum-theoretic predictions were indeed obeyed.



A Gedankenexperiment




A Gedankenexperiment



12GR	32RG	32GR	<u>22RR</u>	<u>33RR</u>	12RG	<u>22GG</u>	13RG	21RG
<u>22GG</u>	<u>11GG</u>	31RG	<u>22RR</u>	21RG	21GG	21GR	<u>11GG</u>	13GR
<u>11GG</u>	31RG	32RG	21GR	31GR	31RG	12GR	23RG	<u>11RR</u>
21GR	13GR	13RG	<u>33GG</u>	21RR	12RG	<u>33RR</u>	23RG	23GR
<u>33GG</u>	<u>33RR</u>	12RR	21RG	31GR	32GR	23RR	13GG	23GG
<u>33RR</u>	12GR	12RG	31GG	<u>33GG</u>	13GG	32GR	32RR	23RG
<u>11RR</u>	31GR	<u>22RR</u>	<u>11RR</u>	23GR	23GR	<u>22GG</u>	13RR	<u>33GG</u>
<u>22GG</u>	<u>11GG</u>	12GR	32RG	32GG	31RR	13RG	<u>22RR</u>	<u>11RR</u>


A Gedankenexperiment



12GR	32RG	32GR	<u>22RR</u>	<u>33RR</u>	12RG	<u>22GG</u>	13RG	21RG
<u>22GG</u>	<u>11GG</u>	31RG	<u>22RR</u>	21RG	21GG	21GR	<u>11GG</u>	13GR
<u>11GG</u>	31RG	32RG	21GR	31GR	31RG	12GR	23RG	<u>11RR</u>
21GR	13GR	13RG	<u>33GG</u>	21RR	12RG	<u>33RR</u>	23RG	23GR
<u>33GG</u>	<u>33RR</u>	12RR	21RG	31GR	32GR	23RR	13GG	23GG
<u>33RR</u>	12GR	12RG	31GG	<u>33GG</u>	13GG	32GR	32RR	23RG
<u>11RR</u>	31GR	<u>22RR</u>	<u>11RR</u>	23GR	23GR	<u>22GG</u>	13RR	<u>33GG</u>
<u>22GG</u>	<u>11GG</u>	12GR	32RG	32GG	31RR	13RG	<u>22RR</u>	<u>11RR</u>

- (1) **Correlation.** If same detector position, then flash same color
- (2) **Uniformity.** If we look only flashes, results are uniform random


A Gedankenexperiment



- (1) **Correlation.** If same detector position, then flash same color
- (2) **Uniformity.** If we look only flashes, results are uniform random

Hypothesis. What is measured by detectors is an intrinsic property of the particles that is fixed before measurement. (Fixed at the creation in C , or along their travel to A or B , but before fixing selecting positions.)


A Gedankenexperiment



Hypothesis. What is measured by detectors is an intrinsic property of the particles that is fixed before measurement.

For type RGR, same flashes for positions 11, 22, 33 or 12, 21.
(Different flashes for positions 12, 31, 23, 32.) Since positions are selected indep. uniformly at random ($1/9$ each pair), so we should obtain same flashes $5/9 > 55\%$ of the time.

A Gedankenexperiment




Hypothesis. What is measured by detectors is an intrinsic property of the particles that is fixed before measurement.

For type RGR, same flashes for positions 11, 22, 33 or 12, 21. (Different flashes for positions 12, 31, 23, 32.) Since positions are selected indep. uniformly at random ($1/9$ each pair), so we should obtain same flashes $5/9 > 55\%$ of the time.

More generally, for types with twice the same letters (RGR, RRG, RGR, GRR, GGR, GRG and RGG) we should observe same flashes 55% of the time. For types RRR and GGG, it should occur 100% of the time.

A Gedankenexperiment



Hypothesis. What is measured by detectors is an intrinsic property of the particles that is fixed before measurement.


More generally, for types with twice the same letters (RGR, RRG, RGR, GRR, GGR, GRG and RGG) we should observe same flashes 55% of the time. For types RRR and GGG, it should occur 100% of the time.

Whatever the type distribution, detectors should flash the same color at least 55% of the time. A contradiction with characteristic (2). Properties do not exist before detection.

The CHSH Game: $a \oplus b = x \wedge y$

Input: $x, y \in \{0, 1\}$

Output: $a, b \in \{0, 1\}$



The CHSH Game: $a \oplus b = x \wedge y$

(1) No deterministic strategies allow Alice & Bob to win.

$$a(0) \oplus b(0) = 0$$

$$a(0) \oplus b(1) = 0$$

$$a(1) \oplus b(0) = 0$$

$$a(1) \oplus b(1) = 1$$

The CHSH Game: $a \oplus b = x \wedge y$

(2) If Alice & Bob share random bits, then they have a strategy to win 75% of the times, and no probabilistic algorithm can do better.

x	y	$x \wedge y$	0	1	x	$\neg x$	y	$\neg y$	$x \oplus y$	$\neg x \oplus y$	success probability
0	0	0	p_1	0	p_3	0	p_5	0	p_7	0	$P_{00} = p_1 + p_3 + p_5 + p_7$
0	1	0	p_1	0	0	p_4	p_5	0	0	p_8	$P_{01} = p_1 + p_4 + p_5 + p_8$
1	0	0	p_1	0	p_3	0	0	p_6	0	p_8	$P_{10} = p_1 + p_3 + p_6 + p_8$
1	1	1	0	p_2	p_3	0	p_5	0	0	p_8	$P_{11} = p_2 + p_3 + p_5 + p_8$

$$\min \{P_{00}, P_{01}, P_{10}, P_{11}\} \leq 3/4$$

because $S = P_{00} + P_{01} + P_{10} + P_{11} =$

$$3p_1 + p_2 + 3p_3 + p_4 + 3p_5 + p_6 + p_7 + p_8 \leq 3 \cdot (p_1 + \dots + p_8) \leq 3.$$

The CHSH Game: $a \oplus b = x \wedge y$

(3) If Alice & Bob have qubits, then they have a strategy to win $\cos^2(\pi/8) > 85\%$ of the times, and no quantum algorithm can do better.

$$|A, B\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} - \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

Alice applies a rotation $R(\theta)$ with $\theta = -\pi/16$ if $x = 0$ and $3\pi/16$ if $x = 1$. Similarly, Bob applies $R(-\pi/16)$ or $R(3\pi/16)$ depending whether $y = 0$ or 1.

The CHSH Game: $a \oplus b = x \wedge y$

- (1) No deterministic strategies allow Alice & Bob to win.
- (2) If Alice & Bob share random bits, then they have a strategy to win 75% of the times, and no probabilistic algorithm can do better.
- (3) If Alice & Bob have qubits, then they have a strategy to win $\cos^2(\pi/8) > 85\%$ of the times, and no quantum algorithm can do better.

The φ -LOCAL Model

(the ultimate non-signalling local model?)



Aim

- + A weaker model than the quantum extension of the LOCAL model for a better understanding of the absolute limits of any conceivable extensions of the LOCAL model by any non-signalling physical theory (preserving causality principle).

Aim

- + A weaker model than the quantum extension of the LOCAL model for a better understanding of the absolute limits of any conceivable extensions of the LOCAL model by any non-signalling physical theory (preserving causality principle).
- + A simple combinatorial model avoiding linear algebra in Hilbert space.

Aim

- + A weaker model than the quantum extension of the LOCAL model for a better understanding of the absolute limits of any conceivable extensions of the LOCAL model by any non-signalling physical theory (preserving causality principle).
- + A simple combinatorial model avoiding linear algebra in Hilbert space.
- Does not give concrete distributed algorithms because instructions with new physics may do not exist yet.

The 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 nodes was affected by input data of $G_{\vec{x}}$ initially localized outside its view, which is

$$\text{view}_t(G_{\vec{x}}, S) := \bigcup_{v \in S} \text{view}_t(G_{\vec{x}}, v)$$

In the φ -LOCAL model, we allow any kind of correlation (possibly non-local) in the output distribution that does not violate locality.


Non-Signalling Distribution

If A, B are nodes at distance more than t , then for an entry (x, y) , the output (a, b) must verify the **non-signalling** property:

$$\Pr(a|x, y) =_{\text{def}} \sum_b \Pr(a, b|x, y) = \Pr(a|x)$$


$$\Pr(b|x, y) =_{\text{def}} \sum_a \Pr(a, b|x, y) = \Pr(b|y) .$$

Non-Signalling Distribution



x	y	a	b		a	b	
0	0	0	0	or	1	1	(proba. 1/2)
0	1	0	0	or	1	1	(proba. 1/2)
1	0	0	0	or	1	1	(proba. 1/2)
1	1	0	1	or	1	0	(proba. 1/2)

Non-Signalling Distribution



x	y	a	b		a	b	
0	0	0	0	or	1	1	(proba. 1/2)
0	1	0	0	or	1	1	(proba. 1/2)
1	0	0	0	or	1	1	(proba. 1/2)
1	1	0	1	or	1	0	(proba. 1/2)

NB: $a \oplus b = x \wedge y$

(4) There is a distribution in the φ -LOCAL model with $t = 0$ round that allows Alice & Bob to win the CHSH game 100% of the time.

View-Based Lower Bounds

The best currently known round lower bounds for many distributed problems, such as MIS, cannot be broken by applying quantum processing, in any conceivable way.

View-Based Lower Bounds

The best currently known round lower bounds for many distributed problems, such as MIS, cannot be broken by applying quantum processing, in any conceivable way.

E.g., MIS requires $\Omega(\sqrt{\log n / \log \log n})$ rounds in general graphs.

View-Based Lower Bounds

The best currently known round lower bounds for many distributed problems, such as MIS, cannot be broken by applying quantum processing, in any conceivable way.

E.g., MIS requires $\Omega(\sqrt{\log n / \log \log n})$ rounds in general graphs.

The $\Omega(\log^* n)$ lower bound for 3-coloring a ring (or MIS in a path) is NOT based on a view argument.

Some Results for φ -LOCAL

- 2-coloring a ring requires $\lceil n/4 \rceil$ rounds, and this is tight.

Some Results for φ -LOCAL

- 2-coloring a ring requires $\lceil n/4 \rceil$ rounds, and this is tight.
- 2-coloring a path requires $\lceil n/3 \rceil$ rounds, and this is tight.

Some Results for φ -LOCAL

- 2-coloring a ring requires $\lceil n/4 \rceil$ rounds, and this is tight.
- 2-coloring a path requires $\lceil n/3 \rceil$ rounds, and this is tight.
- Computing an MIS for a path requires at least two rounds, and three rounds are sufficient.

Some Results for φ -LOCAL

- 2-coloring a ring requires $\lceil n/4 \rceil$ rounds, and this is tight.
- 2-coloring a path requires $\lceil n/3 \rceil$ rounds, and this is tight.
- Computing an MIS for a path requires at least two rounds, and three rounds are sufficient.
- Holroyd-Ligget'16 (stronger model) have constructed a stationary 1-dependent distribution for 4-coloring a path, and that such distribution for 3-coloring is not possible. It follows that 4-coloring can be done in one 1-directional round, and 3-coloring in one (2-directional) round.

Some Results for φ -LOCAL

- 2-coloring a ring requires $\lceil n/4 \rceil$ rounds, and this is tight.
- 2-coloring a path requires $\lceil n/3 \rceil$ rounds, and this is tight.
- Computing an MIS for a path requires at least two rounds, and three rounds are sufficient.
- Holroyd-Ligget'16 (stronger model) have constructed a stationary 1-dependent distribution for 4-coloring a path, and that such distribution for 3-coloring is not possible. It follows that 4-coloring can be done in one 1-directional round, and 3-coloring in one (2-directional) round.
- 3-coloring a path cannot be done in one 1-directional round.

Thank you!