What Can Be Observed Locally?*

Round-Based Models for Quantum Distributed Computing

Cyril Gavoille[1], Adrian Kosowski[1,2], and Marcin Markiewicz[3]

[1] LaBRI - University of Bordeaux
[2] Dept of Algorithms and System Modeling, Gdańsk University of Technology
[3] Institute of Theoretical Physics and Astrophysics, University of Gdańsk

Abstract. We consider the question of *locality* in distributed computing in the context of quantum information. Specifically, we focus on the round complexity of quantum distributed algorithms, with no bounds imposed on local computational power or on the bit size of messages. Linial's $\mathcal{LOCAL}$ model of a distributed system is augmented through two types of quantum extensions: (1) initialization of the system in a quantum entangled state, and/or (2) application of quantum communication channels. For both types of extensions, we discuss proof-of-concept examples of distributed problems whose round complexity is in fact reduced through genuinely quantum effects. Nevertheless, we show that even such quantum variants of the $\mathcal{LOCAL}$ model have non-trivial limitations, captured by a very simple (purely probabilistic) notion which we call "physical locality" (φ-$\mathcal{LOCAL}$). While this is strictly weaker than the "computational locality" of the classical $\mathcal{LOCAL}$ model, it nevertheless leads to a generic view-based analysis technique for constructing lower bounds on round complexity. It turns out that the best currently known lower time bounds for many distributed combinatorial optimization problems, such as *Maximal Independent Set*, bounds cannot be broken by applying quantum processing, in any conceivable way.

1 Introduction

The introduction of computational models based on quantum computing, starting from the works of Deutsch in the 1980's [11], has led to the advent of a new branch of complexity theory. Many studies have for instance focused on the complexity class BQP of problems solvable on a quantum computer in polynomial time with bounded error probability, and its relation to the classical complexity classes. One of the best known algorithmic results in this respect is Shor's polynomial-time method of integer factorization [37] based on the Quantum Fourier Transform, which has recently been partially tested in an experimental set-up for very small values of problem input. Nevertheless, application of quantum information in centralized computing scenarios still proves extremely costly and is riddled with technological difficulties resulting from quantum decoherence

* Supported by the ANR project "ALADDIN", by the INRIA équipe-project "CÉPAGE", and by the KBN Grant 4 T11C 047 25.

I. Keidar (Ed.): DISC 2009, LNCS 5805, pp. 243–257, 2009.

effects. On the other hand, in an even wider time-frame, properties of quantum-mechanical systems have proven to be of interest from the perspective of game theory [4,13,2], information theory [31,22,3], and distributed systems [4,9].

A major line of study (which we briefly look at in the related work section) concerns the application of quantum effects to reduce communication complexity, i.e., to decrease the number of communication bits required to solve a specific task performed within a system graph with several distributed agents. The influence of quantum information on the computing power of distributed systems with node anonymity and distributed systems in the presence of faults has also been studied.

This paper focuses on a different aspect of quantum distributed computing: we do not impose any bounds on the size of communicated messages, but assume that the system operates in synchronous rounds, and ask to what extent quantum effects can reduce the number of rounds required to solve combinatorial optimisation problems. The starting point for considerations is the well-established $\mathcal{LOCAL}$ model a.k.a. Linial's Free model [25,26]. We provide a comparison of the "computational power" of the quantum and non-quantum models, formalising the notion of locality in quantum distributed computing, and showing how it essentially differs from the understanding of locality in the $\mathcal{LOCAL}$ model.

1.1 Related Work

One of the most intensively studied problems related to multi-agent quantum scenarios, when expressed in the language of distributed computing, is roughly trying to address the question: *Can quantum effects be used to enhance distributed computations with messages of bounded size, i.e., in settings inspired by the* $\mathcal{CONGEST}$ *distributed model?* (See [35] for an introduction to the $\mathcal{CONGEST}$ model.) The quantum variant of $\mathcal{CONGEST}$, widely studied in physics, is known as the Local Operations and Classical Communication ($\mathcal{LOCC}$) model. It exploits the key quantum-mechanical concept of an entangled state (see e.g. [31]). This is achieved by altering the initialization phase of the system to allow for a starting state entangled among all the processors, which are locally given quantum computation capabilities; however, communication between processors is still restricted to the exchange of classical information, only. This application of pre-entanglement has been shown to decrease the number of communication bits required to solve certain distributed problems with output collected from one node, and consequently, to decrease the number of required communication rounds when message sizes are bounded. The first proof-of-concept example was provided in [6], where the computation of a specific function of input data distributed among three parties was shown to require at least 3 communicated bits in the classical case, but only 2 communicated bits if the system is initialized in a specific quantum entangled state. Many related results and refinements of this scenario are surveyed in e.g. [7,39].

Other works on the subject have focused on characterising the physical evolution of states attainable in the $\mathcal{LOCC}$ model [30,32,8], while other authors have dealt with the combinatorial complexity of distributing the entangled state

over the whole system in the initialization phase [38]. Other modifications of the model attempt to show that a denser coding of information in transmitted messages is possible when using quantum channels, as compared to classical communication links (see e.g. [5,36]).

Very recently, some authors have begun to study the impact of quantum effects on fundamental concepts of the theory of distributed computing. An overview of this line of research is contained in the recent survey paper [9]. The advantages of applying quantum communication in games against a dynamic adversary are displayed in [1], where it is shown that a constant number of computational rounds is sufficient to solve the quantum Byzantine agreement problem for an n-node system with less than $n/3$ faulty nodes in such a dynamic setting; corresponding classical algorithms require $\Omega(\sqrt{n})$ rounds. Another especially interesting result is that the leader election problem can be solved in distributed systems with quantum links, but no pre-entanglement [40,23]. Some authors have also claimed that problems related to leader election [33,12] and distributed consensus [12,21] can be solved in distributed systems aided by quantum pre-entanglement.

1.2 Outline of the Paper

In Subsection 1.3 we briefly outline the $\mathcal{LOCAL}$ model and its extensions, obtained by modifying the initialization of the system set-up and/or adding quantum communication capabilities on the edges. Whereas this discussion is self-contained, we also provide a formal mathematical definition of the corresponding notions in an extended version of the paper [19]. Subsection 1.4 introduces some notation used when comparing computational models.

In Section 2 we compare the computational power of models based on the proposed extensions of $\mathcal{LOCAL}$. In particular, we prove that adding quantum extensions to the $\mathcal{LOCAL}$ model decreases the round complexity of certain distributed problems. This is achieved through simple proof-of-concept examples.

Most importantly, in Section 3 we introduce a probabilistic framework for proving lower bounds on the distributed time complexity of computational problems in any quantum (or other unconventional) models based on $\mathcal{LOCAL}$. This is directly applied to obtain such lower bounds for many combinatorial optimization problems, including Maximal Independent Set, Greedy Graph Coloring, and problems of spanner construction. As a side effect, the simple concept of "physical locality" formulated in this section, leads to the definition of a computational model we call φ-$\mathcal{LOCAL}$, which appears to be of independent interest.

Finally, in Section 4 we make an attempt to clarify issues with some of the related work on quantum distributed computing as surveyed by [9]. Making use of the framework of computational models defined in the previous sections, we explain why certain claims, saying that problems such as Leader Election or Distributed Consensus benefit from the application of quantum processing, should be approached with caution.

Section 5 contains some concluding remarks and suggests directions of future studies.

1.3 Preliminaries: Description of Computation Models

The $\mathcal{LOCAL}$ Model. The $\mathcal{LOCAL}$ model has been the subject of intensive study in the last 20 years, starting from the seminal works [25,29]. It is assumed that the distributed system consists of a set of processors V (with $|V| = n$) and operates in a sequence of synchronous rounds, each of which involves unbounded computations on the local state variables of the processors, and a subsequent exchange of messages of arbitrary size between pairs of processors which are connected by links (except for round 0, which involves local computations, only). Nodes can identify their neighbours using integer labels assigned successively to communication ports. The local computation procedures encoded in all processors are necessarily the same, and initially all local state variables have the same value for all processors, except for one distinguished local variable $x(v)$ of each processor v which encodes input data. The input of a problem is defined in the form of a labeled graph G_x, where $G = (V, E)$ is the system graph, while $x : V \to \mathbb{N}$ is an assignment of labels to processors. The output of the algorithm is given in the form of a vector of local variables $y : V \to \mathbb{N}$, and the algorithm is assumed to terminate once all variables $y(v)$ are definitely fixed. Herein we assume that faults do not appear on processors and links, that local computation procedures may be randomized (with processors having access to their own generators of random variables), and that the input labels x need not in general be distinct for all processors.

In our considerations, it is convenient to assume that the set of processors V is given *before* the input is defined. This is used for convenience of notation, and neither affects the model, nor the anonymity of nodes in the considered problems.

Extensions of System Initialization ($^{+}\mathcal{S}$ and $^{+}\mathcal{E}$). In the $\mathcal{LOCAL}$ model, it is assumed that the initial set-up of all the processors is identical. This assumption can be relaxed by allowing the processors to obtain some information from a central helper, but only before the start of the distributed process (i.e., independently of the input G_x). The initialization procedure is an integral part of the algorithm used for solving the distributed problem. Several different forms of initialization can be naturally defined; for clarity of discussion, we consider only two extensions of the model: the $^{+}\mathcal{S}$ extension (for *Separable* state), which allows for the most general form of initialization possible in a classical computational setting, and the more powerful $^{+}\mathcal{E}$ extension (for *Entangled* state), which allows for the most general form of initialization available in a quantum distributed system.

The $^{+}\mathcal{S}$ extension. We say that a computational model is equipped with the $^{+}\mathcal{S}$ *extension* if the following modifications are introduced:

- For any computational problem, the computational procedure consists of the distributed algorithm applied by all the processors during the rounds of computation, and an additional (randomized) procedure executed in a centralized way in the initialization phase. The result of the initialization

procedure is an assignment $h : V \to \mathbb{N}$ of *helper* variables to the set of processors. The helper variables are independent[1] of the input G_x.
– For each processor $v \in V$, at the start of round 0, its input label $x(v)$ is augmented by the value $h(v)$, stored in a helper register of the local memory.

It is straightforward to show that the above formulation has two equivalent characterizations. From a computational perspective, we may equivalently say that for each processor v, the helper initialization value $h(v)$ encodes: (1) a unique identifier of v from the range $\{1, \ldots, n\}$, (2) the value of n, (3) the value of a random number, chosen from an arbitrarily large range, and shared by all processors. All further helper information is unnecessary, since it can be computed by the processors in round 0 of the distributed computations (simulation of the centralized assignment of further helper information can be simulated based on random bits and starting information which is common to all processors).

Alternatively, we may say that through the randomized initialization, according to some probability distribution we choose some deterministic initialization of the set of states of individual processors. This intuition precisely corresponds to the notion of a state with uncertainty in classical statistical physics, referred to in quantum-mechanical discussions as a (mixed) *separable state* of the system. It is obviously true to say that *whenever a problem is solved in a model with the* $^{+}\mathcal{S}$ *extension, it may benefit solely from the modification of the system initialization, and not from the laws of quantum mechanics.*

The $^{+}\mathcal{E}$ *extension.* Unlike in classical physics, in quantum mechanics not every initialization of the system has to follow the above pattern. Consider a scenario in which we centrally create an initial global state of the whole system of processors, and spatially distribute "parts" of it to the individual processors (for example, by sharing out among the nodes a set of quantum-correlated photons). Then, each of the processors can perform operations on the "part" of the state assigned to its spatial location; by a loose analogy to processing of classical information, this is sometimes referred to as each processor "manipulating its own quantum bits (qubits)". Given a general initial state of the system, the outcome of such a physical process, as determined by the processors, may display correlations which cannot be described using any classical probabilistic framework. Initial states which can be lead to display such properties are called non-separable, or *entangled states*. Quantum entanglement is without doubt one of the predominant topics studied in quantum-mechanical literature of the last decades; we refer the interested reader to e.g. [31] for an extensive introduction to the topic.

We say that a computational model is equipped with the $^{+}\mathcal{E}$ *extension* if all processors are equipped with helper quantum information registers h, and the computational procedure used to solve a problem sets in the initialization phase in a centralized way some chosen, possibly entangled, quantum state over the set of quantum information registers h of all processors, in a way independent of the input graph G_x.

[1] Helper variables that do depend on the inputs are referred to in the literature as *Oracles* [16,15]. Such extensions are not discussed in this paper.

Extension of Communication Capabilities ($^{+}\mathcal{Q}$). Whereas the application of local quantum operations in each processor does not increase the power of the $\mathcal{LOCAL}$ model as such, the situation changes when the processors can interact with each other using quantum communication channels. Intuitively, such channels allow for the distribution of an entangled state by a processor over several of its neighbours in one communication round; such an effect cannot be achieved using classical communication links.

We say that a computational model is equipped with the $^{+}\mathcal{Q}$ *extension* if all communication links between processors in the system graph are replaced by quantum communication channels.

Models with the Extensions. Modifications to the initialization and communication capabilities of the system are completely independent of each other. For initialization, we can apply no extension, use a separable state ($^{+}\mathcal{S}$), or an entangled state ($^{+}\mathcal{E}$). For communication, we can apply no extension, or use quantum channels ($^{+}\mathcal{Q}$). Hence, we obtain 6 possible models ($\mathcal{LOCAL}$, $\mathcal{LOCAL}^{+}\mathcal{S}$, $\mathcal{LOCAL}^{+}\mathcal{E}$, $\mathcal{LOCAL}^{+}\mathcal{Q}$, $\mathcal{LOCAL}^{+}\mathcal{Q}^{+}\mathcal{S}$, $\mathcal{LOCAL}^{+}\mathcal{Q}^{+}\mathcal{E}$), which are discussed in the following section. (Some of these models collapse onto each other.)

1.4 Notation for Comparing the Power of Computational Models

In order to compare the computational power of different models, we introduce two basic notions: that of the *problem* being solved, and of an *outcome* of the computational process.

Definition 1. *A* problem $\mathscr{P}$ *is a mapping* $G_x \mapsto \{y^i\}$*, which assigns to each input graph* G_x *a set of permissable output vectors* $y^i : V \to \mathbb{N}$*.*

Instead of explicitly saying that we are interested in finding efficient (possibly randomized) distributed algorithms for solving problems within the considered computational models, we characterize the behavior of such procedures through the probability distribution of output vectors which they may lead to, known as an *outcome*. In fact, such a probability distribution is necessarily well defined, whereas formally describing the computational process may be difficult in some unconventional settings (see e.g. the φ-$\mathcal{LOCAL}$ model in Section 3).

Definition 2. *An* outcome $\mathscr{O}$ *is a mapping* $G_x \mapsto \{(y^i, p^i)\}$*, which assigns to each input graph* G_x *a normalized discrete probability distribution* $\{p^i\}$*, such that:* $\forall_i\ p^i > 0$ *and* $\sum_i p^i = 1$*, with* p^i *representing the probability of obtaining* $y^i : V \to \mathbb{N}$ *as the output vector of the distributed system.*

Definition 3. *For any outcome* $\mathscr{O}$ *in a computational model* $\mathcal{M}$ *which is a variant of* $\mathcal{LOCAL}$*, we will write* $\mathscr{O} \in \mathcal{M}[t]$ *if within model* $\mathcal{M}$ *there exists a distributed procedure which yields outcome* $\mathscr{O}$ *after at most* t *rounds of computation.*

We will say that an outcome $\mathscr{O}$ is a *solution* to problem $\mathscr{P}$ *with probability* p if for all G_x, we have: $\sum_{\{(y^i,p^i)\in\mathscr{O}(G_x)\ :\ y^i\in\mathscr{P}(G_x)\}} p_i \geq p$. When $p = 1$, we will simply call $\mathscr{O}$ a *solution* to $\mathscr{P}$ *(with certainty).*

By a slight abuse of notation, for a problem $\mathscr{P}$ we will write $\mathscr{P} \in \mathcal{M}[t]$ (respectively, $\mathscr{P} \in \mathcal{M}[t,p]$) if there exists an outcome $\mathscr{O} \in \mathcal{M}[t]$ which is a solution to problem $\mathscr{P}$ (respectively, a solution to problem $\mathscr{P}$ with probability p).

For two computational models $\mathcal{M}_1$, $\mathcal{M}_2$, we say that $\mathcal{M}_1$ is *not more powerful than* $\mathcal{M}_2$ (denoted $\mathcal{M}_1 \subseteq \mathcal{M}_2$) if for every problem $\mathscr{P}$, for all $t \in \mathbb{N}$ and $p > 0$, $\mathscr{P} \in \mathcal{M}_1[t,p] \implies \mathscr{P} \in \mathcal{M}_2[t,p]$. The relation $\subseteq$ induces a partial order of models which is naturally extended to say that $\mathcal{M}_1$ and $\mathcal{M}_2$ are *equivalent* ($\mathcal{M}_1 = \mathcal{M}_2$), or that $\mathcal{M}_1$ is *less powerful than* $\mathcal{M}_2$ ($\mathcal{M}_1 \subsetneq \mathcal{M}_2$).

It can easily be proved that $\mathcal{M}_1 \subseteq \mathcal{M}_2$ if and only if for every outcome $\mathscr{O}$, for all $t \in \mathbb{N}$, $\mathscr{O} \in \mathcal{M}_1[t] \implies \mathscr{O} \in \mathcal{M}_2[t]$. Such an outcome-based characterisation of models is occasionally more intuitive, since it is not explicitly parameterised by probability p.

In all further considerations, when proving that $\mathcal{M}_1 \subsetneq \mathcal{M}_2$, we will do so in a stronger, deterministic sense, by showing that there exist a problem $\mathscr{P}$ and $t \in \mathbb{N}$ such that $\mathscr{P} \in \mathcal{M}_2[t]$ and $\mathscr{P} \notin \mathcal{M}_1[t]$.

2 Hierarchy of Quantum Models

The most natural variants of $\mathcal{LOCAL}$ which are based on the extensions proposed in the previous subsection are the classical model with separable initialization ($\mathcal{LOCAL}^+\mathcal{S}$), and quantum models with pre-entanglement at initialization, quantum channels, or both ($\mathcal{LOCAL}^+\mathcal{E}$, $\mathcal{LOCAL}^+\mathcal{Q}$, and $\mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E}$, respectively). The strengths of the models can obviously be ordered as follows: $\mathcal{LOCAL} \subseteq \mathcal{LOCAL}^+\mathcal{Q} \subseteq \mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{S} \subseteq \mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E}$, and $\mathcal{LOCAL} \subseteq \mathcal{LOCAL}^+\mathcal{S} \subseteq \mathcal{LOCAL}^+\mathcal{E} \subseteq \mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E}$. We now proceed to show that, whereas $\mathcal{LOCAL}^+\mathcal{E} = \mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E}$, all the remaining inclusions are in fact strict. The hierarchy of the most important models is shown in Fig. 1.

Proposition 1. *$\mathcal{LOCAL} \subsetneq \mathcal{LOCAL}^+\mathcal{S}$. Moreover, there exists a problem $\mathscr{P}$ such that $\mathscr{P} \in \mathcal{LOCAL}^+\mathcal{S}[0]$ and $\mathscr{P} \notin \mathcal{LOCAL}[t]$ for all $t \in \mathbb{N}$.*

Proof. Any problem, which can be solved when given unique node identifiers from the range $\{1, \ldots, n\}$ is clearly in $\mathcal{LOCAL}^+\mathcal{S}[0]$. On the other hand, there are many examples of such problems which are not in $\mathcal{LOCAL}$ (or require $\Omega(n)$ rounds assuming that the system graph is connected and node labels are unique), most trivially the problem $\mathscr{P}$ of assigning unique node identifiers from the range $\{1, \ldots, n\}$ to all nodes. □

More interestingly, one can show that $\mathcal{LOCAL}^+\mathcal{S}$ benefits due to the fact that helper variables $h(v)$ can encode a value which is set in a randomized way. Consider as a simple example a problem $\mathscr{P}'$ whose input is a graph $G = (V, E)$, of sufficiently large order n, with input labels of the nodes encoding unique node identifiers $\{1, \ldots, n\}$ and the value of n; moreover, G is restricted to be the complete graph K_n minus exactly one edge. The goal is to select an edge of the graph, i.e., output y must be such that for some two nodes $u, v \in V$, with $\{u, v\} \in E$, we have $y(u) = y(v) = 1$, and for all other $w \in V$ we have $y(w) = 0$. Even

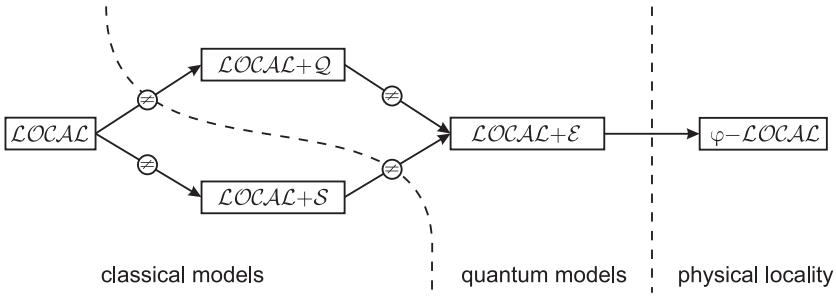

Fig. 1. Hierarchy of computational extensions to the $\mathcal{LOCAL}$ model. See Section 3 for a definition of the φ-$\mathcal{LOCAL}$ model, and Section 1.3 or the extended version [19] for definitions of all other models.

with the knowledge of node identifiers and n, in the $\mathcal{LOCAL}$ model the problem cannot be solved with high probability without communication, i.e., within 0 rounds: we have $\mathscr{P}' \notin \mathcal{LOCAL}[0, e^{-1}]$ (the proof is technical and postponed to the extended version [19]). On the other hand, within the $\mathcal{LOCAL}^{+}\mathcal{S}$ model this problem admits a solution in 0 rounds with probability arbitrarily close to 1 for sufficiently large n. Similar arguments can be applied to display the difference between the models for more advanced problems which simulate collaborative mobile agent scenarios, in particular variants of the cops-and-robbers problems in graphs.

We now point out the difference in power between the classical and quantum models. The proofs proceed by rephrasing one of the best established results of quantum interferometry, first introduced in the context of the so called Bell's Theorem without inequalities, for a 3-particle quantum entangled state (cf. [20] for the original paper or [34] for a contemporary exposition). We use its more algorithmic modulo-4 sum formulation, similar to that found in [41].

Theorem 1. *$\mathcal{LOCAL}^{+}\mathcal{S} \subsetneq \mathcal{LOCAL}^{+}\mathcal{E}$. Moreover, there exists a problem $\mathscr{P}$ such that $\mathscr{P} \in \mathcal{LOCAL}^{+}\mathcal{E}[0]$ and $\mathscr{P} \notin \mathcal{LOCAL}^{+}\mathcal{S}[t]$ for all $t \in \mathbb{N}$.*

Proof. Let $\mathscr{P}$ be a problem defined on a system with 3 nodes. Let the input graph be empty, and assume that input labels $x = (x_1, x_2, x_3) \in \{0,1\}^3$ of respective nodes satisfy the condition $x_1 + x_2 + x_3 \in \{0, 2\}$. An output $y = (y_1, y_2, y_3) \in \{0,1\}^3$ is considered valid for input x if and only if $2(y_1+y_2+y_3) \equiv (x_1+x_2+x_3) \mod 4$. This problem is not in $\mathcal{LOCAL}^{+}\mathcal{S}$, since finding a solution with certainty would imply that there exist three deterministic functions $Y_1, Y_2, Y_3 : \{0,1\} \to \{0,1\}$, such that for any input vector (x_1, x_2, x_3) satisfying the constraints of the problem, $(Y_1(x_1), Y_2(x_2), Y_3(x_3))$ is a valid output vector. It is immediate to show that this is impossible.

The situation is different when the system operates in the $\mathcal{LOCAL}^{+}\mathcal{E}$ model starts in an entangled state. The procedure required to obtain a valid solution is described in detail in [20]. In brief, in the initialization phase we share out to each of the processors one of 3 entangled qubits, carried e.g. by photons,

which are in the entangled tripartite state known as the GHZ state (namely $\frac{1}{\sqrt{2}}(|000\rangle + |111\rangle)$ in Dirac's notation for pure states). Each of the processors then performs a simple transformation on "its own" qubit, in a way dependent only on the processor's input x_i. Finally, a measurement is performed, and it can be shown that the probability distribution of obtained output vectors (the outcome) is that stated in Table 1. Since all of the outputs are accepted as valid for the considered problem $\mathscr{P}$, this implies that $\mathscr{P} \in \mathcal{LOCAL}^{+}\mathcal{E}[0]$. □

Table 1. An outcome $\mathscr{O}$ which is a solution (with certainty) to the modulo-4 sum problem on the 3-node empty graph, and belongs to $\mathcal{LOCAL}^{+}\mathcal{E}[0]$ (see Theorem 1)

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)	1/4	(1, 1, 1)
or (1, 0, 1)	1/4	(1, 0, 0)
or (1, 1, 0)	1/4	(0, 1, 0)
	1/4	(0, 0, 1)

Proposition 2. *$\mathcal{LOCAL} \subsetneq \mathcal{LOCAL}^{+}\mathcal{Q}$. Moreover, for any $t > 0$, there exists a problem $\mathscr{P}$ such that $\mathscr{P} \in \mathcal{LOCAL}^{+}\mathcal{Q}[t]$ and $\mathscr{P} \notin \mathcal{LOCAL}[2t-1]$.*

Proof (sketch). The proof proceeds by a modification of the argument from Theorem 1. This time, we consider a system on $n = 3k+1$ nodes, and an input graph with the topology of a uniformly subdivided star with a central node of degree 3. The modified problem $\mathscr{P}'$ consists in solving the problem from Theorem 1, when the three input and output values are put on the three leaves of the star. Within $\mathcal{LOCAL}$, this problem requires $2k$ rounds to solve, whereas within $\mathcal{LOCAL}^{+}\mathcal{Q}$, k rounds are sufficient. □

Whereas the time distinction between $\mathcal{LOCAL}^{+}\mathcal{S}$ and $\mathcal{LOCAL}^{+}\mathcal{E}$ given by Theorem 1 is remarkable (since it considers the feasibility of solving problems, or when discussing connected graphs, a speed-up from $\Omega(n)$ to 0 communication rounds), the situation is less clear between $\mathcal{LOCAL}^{+}\mathcal{Q}$ and $\mathcal{LOCAL}$. Although a speed-up factor of 2 as expressed by Proposition 2 looks like a natural limit, the authors know of no conclusive arguments to show that it cannot be increased further.

Finally, following the argumentation of [9], we observe that $\mathcal{LOCAL}^{+}\mathcal{E} = \mathcal{LOCAL}^{+}\mathcal{Q}^{+}\mathcal{E}$, or in other words that, given access to pre-entanglement, it is possible to simulate quantum links by means of classical ones. The effect used to achieve this is known as quantum teleportation [34]; by carefully choosing an entangled state over the whole system, it can be applied even when the communicating nodes do not yet know their neighbors' unique identifiers. The amount of pre-entanglement provided at initialization must be sufficient to allow for communication throughout all the rounds of the algorithm.

To complete a discussion of Fig. 1, we point out that $\mathcal{LOCAL}^+\mathcal{Q}$ is incomparable with $\mathcal{LOCAL}^+\mathcal{S}$. This is because the problem discussed in the proof of Proposition 1 belongs to $\mathcal{LOCAL}^+\mathcal{S}$, but not to $\mathcal{LOCAL}^+\mathcal{Q}$, and the problem discussed in the proof of Proposition 2 belongs to $\mathcal{LOCAL}^+\mathcal{Q}[1]$, but not to $\mathcal{LOCAL}^+\mathcal{S}[1]$.

The $\mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{S}$ model has been left out from discussion, since it appears to be of little significance. By considering the same problems as before, we have $\mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{S} \subsetneq \mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E} = \mathcal{LOCAL}^+\mathcal{E}$, so $\mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{S}$ could be placed directly to the left of $\mathcal{LOCAL}^+\mathcal{E}$ in Fig. 1.

3 Lower Time Bounds Based on Physical Locality (φ-$\mathcal{LOCAL}$)

Proving lower bounds on the power of quantum models is problematic. This results, in particular, from the fact that there does not exist as yet an easy-to-use classification of entangled states, or of quantum operations (completely positive maps) which can be performed to transform one quantum state into another. However, in the context of distributed computing, it is possible to consider a more general framework of physical locality, leading to the φ-$\mathcal{LOCAL}$ model we define hereafter, which in turn can be used to bound the power of quantum models.

Within the classical $\mathcal{LOCAL}$ model, we can say that the output of any processor v after t rounds has to be computed based on the input data which can be collected from the input graph G_x by performing an exploration up to a depth of t, starting from node v; we call this the *distance-t local view* denoted by $\mathscr{V}_t(G_x, v)$. This leads to a simple characterisation of the $\mathcal{LOCAL}$ model in terms of valid outcomes (see the extended version [19] for a formalization).

In order to allow for quantum extensions to local, the assumption of classical computability needs to be relaxed, while at the same time retaining in some form the assumption of locality. Given a round-based model with interactions between nearest neighbors only, the physical understanding of locality is as follows: *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* $\mathscr{V}_t(G_x, S) := \bigcup_{v \in S} \mathscr{V}_t(G_x, v)$.

Using the above intuition, we now formalize this notion to obtain what we call the φ-$\mathcal{LOCAL}$ model, i.e., the weakest possible distributed model which still preserves physical locality. Given an output distribution $\{(y^i, p^i)\}$ acting on V, for any subset of vertices $S \subseteq V$ we define its *marginal distribution on set S*, $\{(y^i, p^i)\}[S]$, as the unique distribution $\{(\overline{y}^j, \overline{p}^j)\}$ acting on S which satisfies the condition $\overline{p}^j = \sum_{\{i \,:\, \overline{y}^j = y^i[S]\}} p^i$, where $y^i[S]$ is the restriction of output $y^i : V \to \mathbb{N}$ to nodes from subset $S \subseteq V$.

Definition 4. *An outcome* $G_x \mapsto \{(y^i, p^i)\}$ *belongs to* φ-$\mathcal{LOCAL}[t]$ *if for all subsets* $S \subseteq V$*, for any pair of inputs* $G_x^{(a)}$, $G_x^{(b)}$ *such that* $\mathscr{V}_t(G_x^{(a)}, S) = \mathscr{V}_t(G_x^{(b)}, S)$*, the output distributions corresponding to these inputs have identical marginal distributions on set S, i.e.,* $\{(y^{i(a)}, p^{i(a)})\}[S] = \{(y^{i(b)}, p^{i(b)})\}[S]$.

Quantum relaxations of the $\mathcal{LOCAL}$ model, whether obtained through application of pre-entanglement, quantum channels, or both, lie in terms of strength "in between" the $\mathcal{LOCAL}$ and φ-$\mathcal{LOCAL}$ model. This is expressed by the following theorem, whose proof is provided in the extended version of the paper [19].

Theorem 2. $\mathcal{LOCAL}^{+}\mathcal{Q}^{+}\mathcal{E} \subseteq \varphi\text{-}\mathcal{LOCAL}$.

The theorem captures the property of locality of nearest-neighbor interactions in quantum mechanics, and it does not rely in any way on any other physical concepts, such as causality or speed of information in the theory of relativity.

Although it is not clear whether the containment in the above theorem is strict (we leave this as an open question), the φ-$\mathcal{LOCAL}$ model is still sufficiently constrained to preserve many important lower time bounds known from the $\mathcal{LOCAL}$ model, which are based on arguments of indistinguishability of local views of a node for different inputs. In particular, by careful analysis, it is easy to prove the following statements for the φ-$\mathcal{LOCAL}$ model.

- The problem of finding a maximal independent set in the system graph requires $\Omega(\sqrt{\frac{\log n}{\log\log n}})$ rounds to solve [24].
- The problem of finding a locally minimal (greedy) coloring of the system graph requires $\Omega(\frac{\log n}{\log\log n})$ rounds to solve [18,17].
- The problem of finding a connected subgraph with $O(n^{1+1/k})$ edges requires $\Omega(k)$ rounds to solve [10,14].

The matter is less clear in the case of the $(\Delta+1)$-coloring problem. The proof of the famous lower bound of $\frac{1}{2}\log^* n - O(1)$ rounds [26] (and its extension to randomized algorithms [28]) does not appear to generalize from the $\mathcal{LOCAL}$ model to the φ-$\mathcal{LOCAL}$ model; we are unaware of any (even constant) bound on the number of rounds required to find a solution to $(\Delta+1)$-coloring in φ-$\mathcal{LOCAL}$. Some indication that the technique of coloring neighborhood graphs, used by Linial, may not apply in φ-$\mathcal{LOCAL}$, is that this technique can likewise be used to show a lower bound of $\lfloor \frac{n}{2} \rfloor - 1$ rounds on the time required for 2-coloring the cycle C_n, where n is even. However, in φ-$\mathcal{LOCAL}$ the same problem admits a solution in fewer rounds.

Theorem 3. *The problem of* 2*-coloring the even cycle* C_n *(given unique node labels* x*) belongs to* φ-$\mathcal{LOCAL}[\lceil \frac{n-2}{4} \rceil]$*, but does not belong to* φ-$\mathcal{LOCAL}[\lceil \frac{n-2}{4} \rceil - 1]$.

Proof (sketch). For the lower bound, consider the local view of two nodes u, v which still have disjoint views after $\lceil \frac{n-2}{4} \rceil - 1$ rounds. There are at least two nodes which belong to neither the view of u nor the view of v; hence, u and v cannot distinguish whether they are at an even or at an odd distance from each other in the cycle. This directly leads to the lower bound, since the definition condition of φ-$\mathcal{LOCAL}$ can be shown to be violated for $S = \{u, v\}$.

The upper bound is generated by on outcome $\mathscr{O}$ of the 2-coloring problem, given as follows: each of the 2 legal 2-colorings of C_n is used as the output with probability $\frac{1}{2}$. Such an outcome $\mathscr{O}$ belongs to φ-$\mathcal{LOCAL}[\lceil \frac{n-2}{4} \rceil]$. This can be easily verified, since for any subset $S \subseteq V$ we either have that S consists of

exactly two antipodal nodes of C_n, or the view $\mathscr{V}_{\lceil \frac{n-2}{4} \rceil}(C_{n_x}, S)$ is simply an arc of the cycle. □

It would be interesting to find a constructive quantum procedure for finding a 2-coloring of C_n in $\lceil \frac{n-2}{4} \rceil$ rounds. In particular, we have that 2-coloring of C_6 belongs to φ-$\mathcal{LOCAL}[1]$, does not belong to $\mathcal{LOCAL}^+\mathcal{S}[1]$, and do not know if it belongs to $\mathcal{LOCAL}^+\mathcal{E}[1]$.

4 Simple Problems in a Quantum Setting

In this section, we have a look at some of the related work on quantum distributed problems, as outlined in the survey [9]. Whereas the discussion in this section relies on the results and notation from the preceding sections, it can also be translated into the (not always precisely described) computational models studied in the considered related work.

Two problems which have been used to exhibit the difference between quantum models and non-quantum models are LeaderElection, where the goal is for exactly one node of the system graph to output a value of 1 whereas all other nodes output 0, and a problem which we will call BitPicking, where the goal is for all nodes to return the same output value, either 0 or 1. These discussions include the concept of *fairness*, which in the terminology of this paper means that we are asking not about the problems as such, but about obtaining specific (fair) *outcomes*. More precisely, we will say that FairLeaderElection is the outcome which puts a uniform probability distribution on the n distinct outputs valid for LeaderElection (i.e., on all possible leaders), and FairBitPicking is the outcome which puts a uniform probability distribution on the 2 distinct outputs valid for BitPicking (i.e., picking 0 or 1).

The focus of [33,12,21] is to show that FairBitPicking and FairLeaderElection belong to $\mathcal{LOCAL}^+\mathcal{E}[0]$ (even with some additional restrictions on the amount of allowed pre-entanglement), whereas they do not belong to $\mathcal{LOCAL}[0]$. This statement is correct, however, this effect is due to *the modification of initialization of the system, and not to quantum mechanics*. In fact, we can make the following obvious statement.

Proposition 3. FairBitPicking *and* FairLeaderElection *belong to the non-quantum class* $\mathcal{LOCAL}^+\mathcal{S}[0]$*. Moreover, they can be solved with only one bit of helper information per node, at initialization.*

Finally, we relate to the recent claims that the DistributedConsensus can be solved in a quantum setting without communication. Whereas these claims result from a misunderstanding of the definition [27] of DistributedConsensus, we point out that such a result is impossible in any quantum model, since it is even impossible in φ-$\mathcal{LOCAL}$ (a short proof is provided in the extended version of the paper [19]). We recall that in DistributedConsensus, given an assignment of input labels $(x_1, \ldots, x_n)$ to particular processors, the goal is to obtain an output vector $(y, \ldots, y)$, such that $y \in \{x_1, \ldots, x_n\}$.

Proposition 4. DistributedConsensus $\notin \varphi$-$\mathcal{LOCAL}[0]$.

5 Conclusions and Future Work

We have pointed out that the computational power of quantum variants of the $\mathcal{LOCAL}$ model is *strictly greater* than that of the classical $\mathcal{LOCAL}$ model, or that of the $\mathcal{LOCAL}$ model equipped with helper information such as a pool of shared random bits. It remains to be seen whether a difference can be observed for any problems of practical significance. It is potentially possible that certain combinatorial optimization problems may benefit from quantum extensions to the $\mathcal{LOCAL}$ model. However, we can say that the "view-based" limitations of the $\mathcal{LOCAL}$ model still hold in quantum models. So, one specific question which remains open is whether the $(\Delta+1)$-Coloring problem can be solved in a constant number of rounds in any of the relaxed variants of $\mathcal{LOCAL}$.

Finally, we can ask about a characterization of the limitations of quantum computability, the most natural question being to establish whether the containment $\mathcal{LOCAL}^{+\mathcal{E}} \subseteq \varphi\text{-}\mathcal{LOCAL}$ is strict. As a matter of fact, further studies of the $\varphi\text{-}\mathcal{LOCAL}$ model, which can be seen as the weakest distributed local model, capturing verifiability rather than computability of outcomes, appear to be of interest in their own right.

Acknowledgment. We gratefully thank Pierre Fraigniaud and Zvi Lotker for their preliminary discussions on the EPR effect and its applicability to Distributed Computing. We thank Robert Alicki and Władysław Adam Majewski for helpful discussions concerning quantum dynamic maps, and Marek Żukowski for several references on quantum information.

References

1. Ben-Or, M., Hassidim, A.: Fast quantum byzantine agreement. In: 37^{th} Annual ACM Symposium on Theory of Computing (STOC), pp. 481–485. ACM Press, New York (2005)
2. Benjamin, S.C., Hayden, P.M.: Multiplayer quantum games. Physical Review A 64(3), 030301 (2001)
3. Bennett, C.H., Shor, P.W.: Quantum information theory. IEEE Transactions on Information Theory 44, 2724–2742 (1998)
4. Broadbent, A., Tapp, A.: Can quantum mechanics help distributed computing? ACM SIGACT News - Distributed Computing Column 39(3), 67–76 (2008)
5. Buhrman, H., Cleve, R., Wigderson, A.: Quantum vs. classical communication and computation. In: 30^{th} Annual ACM Symposium on the Theory of Computing (STOC), pp. 63–68 (1998)
6. Cleve, R., Buhrman, H.: Substituting quantum entanglement for communication. Physical Review A 56(2), 1201–1204 (1997)
7. de Wolf, R.: Quantum communication and complexity. Theoretical Computer Science 287(1), 337–353 (2002)
8. den Nest, M.V., Dür, W., Vidal, G., Briegel, H.: Classical simulation versus universality in measurement-based quantum computation. Physical Review A 75(1), 012337 (2007)
9. Denchev, V.S., Pandurangan, G.: Distributed quantum computing: A new frontier in distributed systems or science fiction? ACM SIGACT News - Distributed Computing Column 39(3), 77–95 (2008)

10. Derbel, B., Gavoille, C., Peleg, D., Viennot, L.: On the locality of distributed sparse spanner construction. In: 27^{th} Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 273–282. ACM Press, New York (2008)
11. Deutsch, D.: Quantum theory, the Church-Turing principle and the universal quantum computer. Proceedings of the Royal Society of London A400, 97–117 (1985)
12. D'Hondt, E., Panangaden, P.: The computational power of the W and GHZ states. Quantum Information and Computation 6(2), 173–183 (2006)
13. Eisert, J., Wilkens, M., Lewenstein, M.: Quantum games and quantum strategies. Physical Review Letters 83(11), 3077–3080 (1999)
14. Elkin, M.: A near-optimal fully dynamic distributed algorithm for maintaining sparse spanners. In: 26^{th} Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 195–204. ACM Press, New York (2007)
15. Fraigniaud, P., Gavoille, C., Ilcinkas, D., Pelc, A.: Distributed computing with advice: Information sensitivity of graph coloring. In: Arge, L., Cachin, C., Jurdziński, T., Tarlecki, A. (eds.) ICALP 2007. LNCS, vol. 4596, pp. 231–242. Springer, Heidelberg (2007)
16. Fraigniaud, P., Ilcinkas, D., Pelc, A.: Oracle size: a new measure of difficulty for communication tasks. In: 25^{th} Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 179–187. ACM Press, New York (2006)
17. Gavoille, C., Klasing, R., Kosowski, A., Kuszner, Ł., Navarra, A.: On the complexity of distributed graph coloring with local minimality constraints. Networks (to appear, 2009)
18. Gavoille, C., Klasing, R., Kosowski, A., Navarra, A.: Brief announcement: On the complexity of distributed greedy coloring. In: Pelc, A. (ed.) DISC 2007. LNCS, vol. 4731, pp. 482–484. Springer, Heidelberg (2007)
19. Gavoille, C., Kosowski, A., Markiewicz, M.: What can be observed locally? Round-based models for quantum distributed computing. Technical report, arXiv: quant-ph/0903.1133 (2009)
20. Greenberger, D.M., Horne, M.A., Zeilinger, A.: Going beyond Bell's Theorem. In: Bell's Theorem, Quantum Theory, and Conceptions of the Universe, pp. 69–72. Kluwer, Dordrecht (1989)
21. Helm, L.: Brief announcement: Quantum distributed consensus. In: 27th Annual ACM Symposium on Principles of Distributed Computing (PODC), p. 445. ACM Press, New York (2008)
22. Jaeger, G.: Quantum Information. An Overview. Springer, Heidelberg (2007)
23. Kobayashi, H., Matsumoto, K., Tani, S.: Fast exact quantum leader election on anonymous rings. In: 8^{th} Asian Conference on Quantum Information Science (AQIS), August 2008, pp. 157–158 (2008)
24. Kuhn, F., Moscibroda, T., Wattenhofer, R.: What cannot be computed locally! In. In: 23^{rd} Annual ACM Symposium on Principles of Distributed Computing (PODC), pp. 300–309. ACM Press, New York (2004)
25. Linial, N.: Distributive graph algorithms - Global solutions from local data. In: 28^{th} Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 331–335. IEEE Computer Society Press, Los Alamitos (1987)
26. Linial, N.: Locality in distributed graphs algorithms. SIAM Journal on Computing 21(1), 193–201 (1992)
27. Lynch, N.: Distributed Algorithms. Morgan Kaufmann Publishers, San Francisco (1997)
28. Naor, M.: A lower bound on probabilistic algorithms for distributive ring coloring. SIAM Journal on Discrete Mathematics 4(3), 409–412 (1991)

29. Naor, M., Stockmeyer, L.: What can be computed locally. SIAM Journal on Computing 24(6), 1259–1277 (1995)
30. Nielsen, M.: Conditions for a class of entanglement transformations. Physical Review Letters 83(2), 436–439 (1999)
31. Nielsen, M., Chuang, I.: Quantum Computation and Quantum Information. Cambridge University Press, Cambridge (2000)
32. Owari, M., Matsumoto, K., Murao, M.: Entanglement convertibility for infinite-dimensional pure bipartite states. Physical Review A 70(5), 1–4 (2004)
33. Pal, S.P., Singh, S.K., Kumar, S.: Multi-partite quantum entanglement versus randomization: Fair and unbiased leader election in networks. Technical report, arXiv: quant-ph/0306195v1 (June 2003)
34. Pan, J.-W., Chen, Z.-B., Żukowski, M., Weinfurter, H., Zeilinger, A.: Multi-photon entanglement and interferometry. Technical report, arXiv: quant-ph/0805.2853v1 (May 2008)
35. Peleg, D.: Proximity-preserving labeling schemes and their applications. In: Widmayer, P., Neyer, G., Eidenbenz, S. (eds.) WG 1999. LNCS, vol. 1665, pp. 30–41. Springer, Heidelberg (1999)
36. Raz, R.: Exponential separation of quantum and classical communication complexity. In: 31^{st} Annual ACM Symposium on the Theory of Computing (STOC), pp. 358–367 (1999)
37. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Journal on Computing 26(5), 1484–1509 (1997)
38. Singh, S.K., Kumar, S., Pal, S.P.: Characterizing the combinatorics of distributed EPR pairs for multi-partite entanglement. Technical report, arXiv: quant-ph/0306049v2 (January 2004)
39. Ta-Shma, A.: Classical versus quantum communication complexity. SIGACT News 30(3), 25–34 (1999)
40. Tani, S., Kobayashi, H., Matsumoto, K.: Exact quantum algorithms for the leader election problem. In: Diekert, V., Durand, B. (eds.) STACS 2005. LNCS, vol. 3404, pp. 581–592. Springer, Heidelberg (2005)
41. Żukowski, M.: On Bell's Theorem, quantum communication, and entanglement detection. In: Foundations of Probability and Physics 5 (August 2008)