

What Can be Observed Locally?

Round-based Models for Quantum Distributed Computing

Cyril Gavaille*
LaBRI - University of Bordeaux
Talence, France

Adrian Kosowski*
LaBRI - University of Bordeaux
Talence, France

Marcin Markiewicz
Institute of Theoretical Physics and Astrophysics
University of Gdańsk, Poland

Abstract

In this paper we focus on the question of *locality* in distributed computing in the context of quantum information. To the best of our knowledge, this is the first work which focuses solely on the round complexity of quantum distributed algorithms, with no bounds imposed on local computational power or on the bit size of messages. Specifically, we consider Linial’s \mathcal{LOCAL} model of a distributed system 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 put forward valid 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” ($\varphi\text{-}\mathcal{LOCAL}$). 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. Specifically, 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.

*Supported by the ANR project “ALADDIN”, and the INRIA équipe-project “CÉPAGE”.

1 Introduction

The introduction of computational models based on quantum computing, starting from the works of Deutsch in the 1980's [Deu85], 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 [Sho94, Sho97, Buh96] based on the Quantum Fourier Transform, which has recently been partially tested in an experimental set-up for very small values of problem input [LWL⁺07, CYL07]. Nevertheless, application of quantum information in centralized computing scenarios still proves extremely costly and is riddled with technological difficulties resulting from quantum decoherence effects [Sch05]. 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 [BT08, EWL99, BH01], information theory [NC00, Jae07, BS98], and distributed systems [BT08, DP08].

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 *LOCAL* model a.k.a. Linial's Free model [Lin87, Lin92]. 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 *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 CONGEST distributed model¹?* The quantum variant of *CONGEST*, widely studied in physics, is known as the *LOCC* model². It exploits the key quantum-mechanical concept of an entangled state (see e.g. [NC00, HHH96]). 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 [CB97], where the computation of

¹See [Pel99] for an introduction to the *CONGEST* model.

²It stands for *Local Operations and Classical Communication*.

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. [Z08, BCvD01].

Other works on the subject have focused on characterising the physical evolution of states attainable in the \mathcal{LOCC} model [Nie99, OMM04, dNDVB07], while other authors have dealt with the combinatorial complexity of distributing the entangled state over the whole system in the initialization phase [SKP04]. 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. [BCdWZ99]).

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 by Denchev and Pandurangan [DP08]. The advantages of applying quantum communication in games against a dynamic adversary are displayed in [BOH05], 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 [TKM05, KMT08]. Some authors have also claimed that problems related to leader election [PSK03, DP06] and distributed consensus [DP06, Hel08] 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 Appendix A. 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 [DP08]. 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

In this section we briefly recall the computational properties of the *LOCAL* model, which has been the subject of intensive study in the last 20 years, starting from the seminal works of [Lin87, NS95]. When considering the *LOCAL* model in the context of quantum processing, it has to be noted that simply introducing a “quantum computer” as a module in each processor does not affect the power of the model, since in *LOCAL* the processors as such are already assumed to have unbounded capabilities of local computation.

There exist two distinct and independent approaches to extending the *LOCAL* model: by modifying the initial set-up of the system (leading to extensions which we call ^+S and ^+E), and by introducing quantum communication channels (the ^+Q extension). Of these three extensions, two (^+E and ^+Q) rely on quantum processing and roughly correspond to settings studied in some related work [DP08], whereas the third extension (^+S) is purely computational in the classical sense, and is introduced in this work.

The discussion which follows is intentionally informal, whereas rigorous definitions and some further considerations are postponed to Appendix A. The formalism in the Appendix is used in particular for showing computational limitations of models and pointing out errors in previous work, hence we keep it precise in a mathematical sense and free from any implicit assumptions.

The *LOCAL* model. 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 $x : V \rightarrow \mathbb{N}$ is an assignment of labels to processors. The output of the algorithm is given in the form of a vector of local variables $y : V \rightarrow \mathbb{N}$, 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 unique 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 does not affect neither the model in any way, nor the anonymity of nodes in the considered problems.

Initialization of the system (^+S and ^+E extensions). In the *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 ^+S extension (for *Separable* state), which allows for the most general form of initialization possible in a classical computational setting, and the more powerful ^+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 $h : V \rightarrow \mathbb{N}$ of *helper* variables to the set of processors. The helper variables are independent³ 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 $\{1, \dots, n\}$, (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, coming from a single SPDC⁴ emission process). 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 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. [NC00] for an extensive introduction to the topic.

³Helper variables that do depend on the inputs are referred to in the literature as *Oracles* [FIP06, FGIP07]. Such extensions are not discussed in this paper.

⁴Spontaneous Parametric Down-Conversion.

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 .

Of course, the definition of the ${}^+\mathcal{E}$ extension does not require that the starting state is entangled; for the special case when it is separable, the ${}^+\mathcal{E}$ extension is precisely equivalent to the ${}^+\mathcal{S}$ extension.

Communication capabilities (${}^+\mathcal{Q}$ extension). 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 based on 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 (message exchanges with classical information), 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 collapse onto each other, in particular, $\mathcal{LOCAL}{}^+\mathcal{Q}{}^+\mathcal{E}$ and $\mathcal{LOCAL}{}^+\mathcal{E}$ are equivalent in terms of computational power (Proposition 4).

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* \mathcal{P} is a mapping $G_x \mapsto \{y^i\}$, which assigns to each input graph G_x a set of permissible output vectors $y^i : V \rightarrow \mathbb{N}$.

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* \mathcal{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 $y^i : V \rightarrow \mathbb{N}$ as the output vector of the distributed system.

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

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

By a slight abuse of notation, for a problem \mathcal{P} we will write $\mathcal{P} \in \mathcal{M}[t]$ (respectively, $\mathcal{P} \in \mathcal{M}[t, p]$) if there exists an outcome $\mathcal{O} \in \mathcal{M}[t]$ which is a solution to problem \mathcal{P} (respectively, a solution to problem \mathcal{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 \mathcal{P} , for all $t \in \mathbb{N}$ and $p > 0$, $\mathcal{P} \in \mathcal{M}_1[t, p] \implies \mathcal{P} \in \mathcal{M}_2[t, p]$. 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 \mathcal{O} , for all $t \in \mathbb{N}$, $\mathcal{O} \in \mathcal{M}_1[t] \implies \mathcal{O} \in \mathcal{M}_2[t]$. 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 \mathcal{P} and $t \in \mathbb{N}$ such that $\mathcal{P} \in \mathcal{M}_2[t]$ and $\mathcal{P} \notin \mathcal{M}_1[t]$.

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

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

Proof. Any problem, which can be solved when given unique node identifiers from the range $\{1, \dots, n\}$ is clearly in $\mathcal{LOCAL}^+S[0]$. 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 \mathcal{P} of assigning unique node identifiers from the range $\{1, \dots, n\}$ to all nodes. \square

More interestingly, one can show that \mathcal{LOCAL}^+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 \mathcal{P}' whose input is a graph $G = (V, E)$, of sufficiently large order n , with input labels of the nodes encoding unique node identifiers $\{1, \dots, n\}$ 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

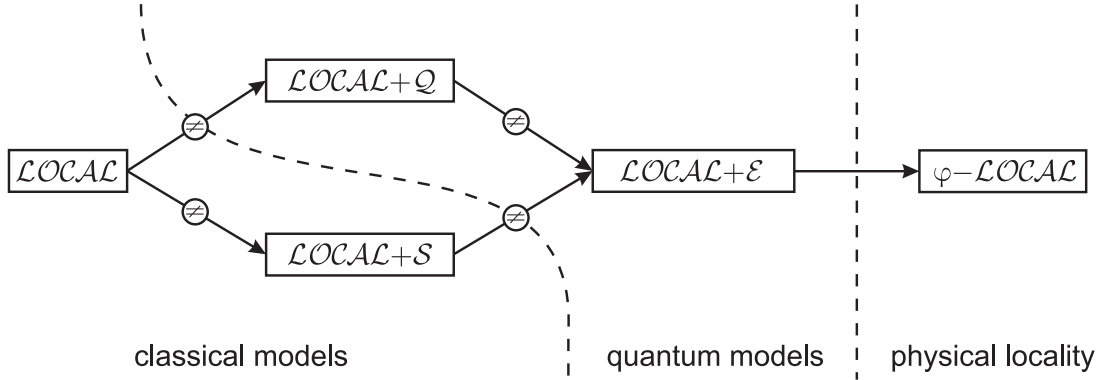


Figure 1: Hierarchy of computational extensions to the \mathcal{LOCAL} model. See Section 3 for a definition of the φ - \mathcal{LOCAL} model, and Section 1.3 or Appendix A for definitions of all other models.

$\{u, v\} \in E$, we have $y(u) = y(v) = 1$, and for all other $w \in V$ we have $y(w) = 0$. Even 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 $\mathcal{P}' \notin \mathcal{LOCAL}[0, e^{-1}]$ (the proof is technical, see Appendix C.1). 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. [GHZ89] for the original paper, [Mer90] for a very informal intuition, or [PCZ⁺08] for a contemporary exposition). We use its more algorithmic modulo-4 sum formulation, similar to that found in [Z08].

Theorem 2. $\mathcal{LOCAL}+\mathcal{S} \subsetneq \mathcal{LOCAL}+\mathcal{E}$. Moreover, there exists a problem \mathcal{P} such that $\mathcal{P} \in \mathcal{LOCAL}+\mathcal{E}[0]$ and $\mathcal{P} \notin \mathcal{LOCAL}+\mathcal{S}[t]$ for all $t \in \mathbb{N}$.

Proof. Let \mathcal{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 $2(y_1 + y_2 + y_3) \equiv (x_1 + x_2 + x_3) \pmod{4}$. This problem is not in $\mathcal{LOCAL}+\mathcal{S}$, since finding a solution with certainty would imply that there exist three deterministic functions $Y_1, Y_2, Y_3 : \{0, 1\} \rightarrow \{0, 1\}$, 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.⁵

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 [GHZ89]. 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

⁵By considering all of the possible inputs, we obtain the set of equations: $2(Y_1(0) + Y_2(0) + Y_3(0)) \equiv 0 \pmod{4}$, and $2(Y_1(1) + Y_2(1) + Y_3(0)) \equiv 2(Y_1(1) + Y_2(0) + Y_3(1)) \equiv 2(Y_1(0) + Y_2(1) + Y_3(1)) \equiv 2 \pmod{4}$, which is contradictory: by summing the left-hand sides of all four equations we obtain $0 \equiv 6 \pmod{4}$.

Table 1: An outcome \mathcal{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 2).

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

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 \mathcal{P} , this implies that $\mathcal{P} \in \mathcal{LOCAL}^{\mathcal{E}}[0]$. \square

We note that the obtained outcome $\mathcal{O} \in \mathcal{LOCAL}^{\mathcal{E}}[0]$ is a solution to \mathcal{P} with certainty, but it is not deterministic, yielding different outputs with probability 1/4 (Table 1); in fact, within $\mathcal{LOCAL}^{\mathcal{E}}$ there does not exist an outcome which is a solution to \mathcal{P} , and yields some output with probability 1. This sort of situation could not occur in \mathcal{LOCAL} , or in any other classical model.

Proposition 3. $\mathcal{LOCAL} \subsetneq \mathcal{LOCAL}^{\mathcal{Q}}$. Moreover, for any $t > 0$, there exists a problem \mathcal{P} such that $\mathcal{P} \in \mathcal{LOCAL}^{\mathcal{Q}}[t]$ and $\mathcal{P} \notin \mathcal{LOCAL}[2t - 1]$.

Proof. The proof proceeds by a modification of the argument from Theorem 2. 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 \mathcal{P}' consists in solving the problem from Theorem 2, 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, since the three leaves are at a distance of $2k$ from each other, and need to communicate to solve the problem. On the other hand, in $\mathcal{LOCAL}^{\mathcal{Q}}$ we are given quantum communication links. Hence, in round 0, the central node can create an entangled tripartite GHZ state, and propagate its qubits in k rounds⁶ to the leaves of the graph, which then apply the previously discussed quantum procedure. \square

Whereas the time distinction between $\mathcal{LOCAL}^{\mathcal{S}}$ and $\mathcal{LOCAL}^{\mathcal{E}}$ given by Theorem 2 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 3 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 [DP08], we note 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

⁶Observe that intermediate nodes simply send the qubit on, without making any copies.

links by means of classical ones. The effect used to achieve this is known as quantum teleportation [PCZ⁺08]; 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.

Proposition 4 ([DP08]). $\mathcal{LOCAL}^{\mathcal{E}} = \mathcal{LOCAL}^{\mathcal{Q}+\mathcal{E}}$.

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 3 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 $\mathcal{V}_t(G_x, v)$. This leads to a simple characterisation of the \mathcal{LOCAL} model in terms of valid outcomes (see Appendix C.2 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, since it is an essential part of physical theory as we understand it today (cf. e.g. [Shi84, Str07] for different approaches to the problem). To define locality, for a moment we choose to look at the system from a physicist's perspective, with the distributed system as an experimental stand, with processors as black boxes, with input data G_x as part of the experimental set-up, and with output y as the data resulting of a single experiment. For each input, the experiment is performed for an ensemble of identical systems, obtaining a probability distribution of outputs $\{(y^i, p^i)\}$. Now, 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 $\mathcal{V}_t(G_x, S) := \bigcup_{v \in S} \mathcal{V}_t(G_x, v)$.*

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 $\{(\bar{y}^j, \bar{p}^j)\}$ acting on S which satisfies the condition $\bar{p}^j = \sum_{\{i : \bar{y}^j = y^i[S]\}} p^i$, where $y^i[S]$ is the restriction of output $y^i : V \rightarrow \mathbb{N}$ 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 $\mathcal{V}_t(G_x^{(a)}, S) = \mathcal{V}_t(G_x^{(b)}, S)$, the output distri-

butions 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 $\varphi\text{-}\mathcal{LOCAL}$ model. This is expressed by the following theorem, whose proof we perform in Appendix B (after previously reformulating the models in equivalent mathematical terms for easier manipulation in Appendix A).

Theorem 5. $\mathcal{LOCAL}^+Q^+\mathcal{E} \subseteq \varphi\text{-}\mathcal{LOCAL}$.

The theorem captures the property of locality of nearest-neighbor interactions in quantum mechanics, and its proof can be seen as a boundary case (for discrete rounds) of the more physical continuous-time setting studied in [BR81]. 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 $\varphi\text{-}\mathcal{LOCAL}$ 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 $\varphi\text{-}\mathcal{LOCAL}$ 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 [KMW04].
- The problem of finding a locally minimal (greedy) coloring of the system graph requires $\Omega(\frac{\log n}{\log \log n})$ rounds to solve [GKKN07, GKK⁺09].
- The problem of finding a connected subgraph with $O(n^{1+1/k})$ edges requires $\Omega(k)$ rounds to solve [DGPV08, Elk07].

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 [Lin92] (and its extension to randomized algorithms [Nao91]) does not appear to generalize from the \mathcal{LOCAL} model to the $\varphi\text{-}\mathcal{LOCAL}$ model; we are unaware of any (even constant) bound on the number of rounds required to find a solution to $(\Delta + 1)$ -coloring in $\varphi\text{-}\mathcal{LOCAL}$. Some indication that the technique of coloring neighborhood graphs, used by Linial, may not apply in $\varphi\text{-}\mathcal{LOCAL}$, 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 $\varphi\text{-}\mathcal{LOCAL}$ the same problem admits a solution in fewer rounds.

Theorem 6. *The problem of 2-coloring the even cycle C_n (given unique node labels x) belongs to $\varphi\text{-}\mathcal{LOCAL}[\lceil \frac{n-2}{4} \rceil]$, but does not belong to $\varphi\text{-}\mathcal{LOCAL}[\lceil \frac{n-2}{4} \rceil - 1]$.*

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 $\varphi\text{-}\mathcal{LOCAL}$ can be shown to be violated for $S = \{u, v\}$.

The upper bound is generated by on outcome \mathcal{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

\mathcal{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 $\mathcal{V}_{\lceil \frac{n-2}{4} \rceil}(C_n, S)$ is simply an arc of the cycle. \square

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 [DP08]. 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 1.⁷ 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 [PSK03, DP06, Hel08] 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 7. *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.*

Proof. There is no input for the considered outcomes, hence the initialization procedure can be defined so as to encode the appropriate output vector in the helper data $h(v)$, choosing specific outputs according to the required probability distribution. \square

Consequently, this sort of study should be considered in the context of the $\mathcal{LOCAL}^+\mathcal{S}$ modification, or in other words, the benefits of adding purely classical helper information to the \mathcal{LOCAL} model (and not what some authors refer to as “quantum non-locality”). Whereas in a formal sense it is not a mistake to say that such an effect can also be obtained when using a quantum entangled state as the “helper”, this is technologically difficult to implement, complicates the discussion, and does not save information in any way, since for the considered outcomes, the required helper data can already be encoded using one bit per node within $\mathcal{LOCAL}^+\mathcal{S}$.

⁷Two sets of authors [DP06, Hel08] confuse the latter question with that of the **DistributedConsensus** problem [Lyn97, AW04].

As such, this sort of approach can be seen as useless from the perspective of distributed computing.⁸ In order to capture the benefit coming from the quantum setup, one has to display quantum correlations which cannot be described in the classical framework (Theorem 2 and Proposition 3).

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 [Lyn97, AW04] of `DistributedConsensus`, we point out that such a result is impossible in any quantum model, since it is even impossible in φ -`LOCAL`. We recall that in `DistributedConsensus`, given an assignment of input labels (x_1, \dots, x_n) to particular processors, the goal is to obtain an output vector (y, \dots, y) , such that $y \in \{x_1, \dots, x_n\}$.

Proposition 8. `DistributedConsensus` \notin φ -`LOCAL`[0].

Proof. Consider a system with only two processors, having inputs $x_1, x_2 \in \{0, 1\}$. Let outcome \mathcal{O} be a valid solution to `DistributedConsensus` with certainty. Then, \mathcal{O} must be given as the following mapping $x \mapsto \{(p^i, y^i)\}$ for some probability values $p, q \in [0, 1]$: $(0, 0) \mapsto \{(1, (0, 0))\}$, $(1, 1) \mapsto \{(1, (1, 1))\}$, $(0, 1) \mapsto \{(p, (0, 0)), (1 - p, (1, 1))\}$, and $(1, 0) \mapsto \{(q, (0, 0)), (1 - q, (1, 1))\}$. Now, suppose that $\mathcal{O} \in \varphi$ -`LOCAL`[0]. Applying the definition of φ -`LOCAL` to set S consisting of processor 1 only, considering inputs $x_a = (0, 0)$ and $x_b = (0, 1)$, we obtain $p = 1$. Likewise, applying the same definition to set S' consisting of processor 2 only, considering inputs $x'_a = (1, 1)$ and $x'_b = (0, 1)$, we obtain $p = 0$, a contradiction. \square

5 Conclusions and Future Work

We have pointed out that the computational power of quantum variants of the `LOCAL` model is *strictly greater* than that of the classical `LOCAL` model, or that of the `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 `LOCAL` model. However, we can say that the “view-based” limitations of the `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 `LOCAL`.

Finally, we can ask about a characterization of the limitations of quantum computability, the most natural question being to establish whether the containment $\text{LOCAL}^{\mathcal{E}} \subseteq \varphi\text{-LOCAL}$ is strict. As a matter of fact, further studies of the φ -`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.

⁸Some observations can perhaps be of purely physical interest. For example, for a problem with empty input, whenever there exists a one-bit helper function $h : V \rightarrow \{0, 1\}$ leading to the desired output in `LOCAL`+ S [0], unique up to negation of 0 and 1, there will also exist an entangled pure n -qubit state spread over the nodes leading to the desired output in `LOCAL`+ \mathcal{E} [0], unique up to transformation of the local basis $\{|0\rangle, |1\rangle\}$. [DP06] note that this is the case for `FairBitPicking` and `FairLeaderElection`.

References

- [AW04] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics (2nd ed.)*. Wiley-Interscience Publication, 2004.
- [Bae87] J. Baez. Bell’s inequality for C*-algebras. *Letters in Mathematical Physics*, 13(2):136–137, February 1987.
- [BCdWZ99] H. Buhrman, R. Cleve, R. de Wolf, and C. Zalka. Bounds for small-error and zero-error quantum algorithms. In *40th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 358–368. IEEE Computer Society Press, October 1999.
- [BCvD01] H. Buhrman, R. Cleve, and W. van Dam. Quantum entanglement and communication complexity. *SIAM Journal on Computing*, 30(6):1829–1841, 2001.
- [BH01] S. C. Benjamin and P. M. Hayden. Multiplayer quantum games. *Physical Review A*, 64(3):030301, 2001.
- [BOH05] M. Ben-Or and A. Hassidim. Fast quantum byzantine agreement. In *37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 481–485. ACM Press, May 2005.
- [BR79] O. Bratelli and D. W. Robinson. *Operator algebras and quantum statistical mechanics*, volume I. Springer-Verlag, 1979.
- [BR81] O. Bratelli and D. W. Robinson. *Operator algebras and quantum statistical mechanics*, volume II. Springer-Verlag, 1981.
- [BS98] C. H. Bennett and P. W. Shor. Quantum information theory. *IEEE Transactions on Information Theory*, 44:2724–2742, 1998.
- [BT08] A. Broadbent and A. Tapp. Can quantum mechanics help distributed computing? *ACM SIGACT News - Distributed Computing Column*, 39(3):67–76, September 2008.
- [Buh96] H. Buhrman. A short note on Shor’s factoring algorithm. *ACM SIGACT News*, 27(1):89–90, March 1996.
- [CB97] R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, Aug 1997.
- [CYL07] T. Y. J.-W. P. Chao-Yang Lu, Daniel E. Browne. Demonstration of shor’s quantum factoring algorithm using photonic qubits. *Physical Review Letters*, 99(25):4, 2007.
- [Deu85] D. Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London*, A400:97–117, 1985.
- [DGPV08] B. Derbel, C. Gavoille, D. Peleg, and L. Viennot. On the locality of distributed sparse spanner construction. In *27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 273–282. ACM Press, August 2008.
- [dNDVB07] M. V. den Nest, W. Dür, G. Vidal, and H. Briegel. Classical simulation versus universality in measurement-based quantum computation. *Physical Review A*, 75(1):012337, 2007.
- [DP06] E. D’Hondt and P. Panangaden. The computational power of the W and GHZ states. *Quantum Information and Computation*, 6(2):173–183, March 2006.
- [DP08] V. S. Denchev and G. Pandurangan. Distributed quantum computing: A new frontier in distributed systems or science fiction? *ACM SIGACT News - Distributed Computing Column*, 39(3):77–95, September 2008.
- [Elk07] M. Elkin. A near-optimal fully dynamic distributed algorithm for maintaining sparse spanners. In *26th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 195–204. ACM Press, August 2007.
- [EWL99] J. Eisert, M. Wilkens, and M. Lewenstein. Quantum games and quantum strategies. *Physical Review Letters*, 83(11):3077, 1999.

- [FGIP07] P. Fraigniaud, C. Gavoille, D. Ilcinkas, and A. Pelc. Distributed computing with advice: Information sensitivity of graph coloring. In *34th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 4596 of Lecture Notes in Computer Science, pages 231–242. Springer, July 2007.
- [FIP06] P. Fraigniaud, D. Ilcinkas, and A. Pelc. Oracle size: a new measure of difficulty for communication tasks. In *25th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 179–187. ACM Press, July 2006.
- [GHZ89] D. M. Greenberger, M. A. Horne, and A. Zeilinger. Going beyond Bell’s Theorem. In *Bell’s Theorem, Quantum Theory, and Conceptions of the Universe*, pages 69–72. Kluwer, 1989.
- [GKK⁺09] C. Gavoille, R. Klasing, A. Kosowski, L. Kuszner, and A. Navarra. On the complexity of distributed graph coloring with local minimality constraints. *Networks*, 2009. To appear.
- [GKKN07] C. Gavoille, R. Klasing, A. Kosowski, and A. Navarra. Brief announcement: On the complexity of distributed greedy coloring. In *21st International Symposium on Distributed Computing (DISC)*, volume 4731 of Lecture Notes in Computer Science, pages 482–484. Springer, September 2007.
- [Hel08] L. Helm. Brief announcement: Quantum distributed consensus. In *27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 445–445. ACM Press, August 2008.
- [HHH96] M. Horodecki, P. Horodecki, and R. Horodecki. Separability of mixed states: necessary and sufficient conditions. *Physics Letters A*, 223:1–8, 1996.
- [Jae07] G. Jaeger. *Quantum Information. An Overview*. Springer-Verlag, 2007.
- [KMT08] H. Kobayashi, K. Matsumoto, and S. Tani. Fast exact quantum leader election on anonymous rings. In *8th Asian Conference on Quantum Information Science (AQIS)*, pages 157–158, August 2008.
- [KMW04] F. Kuhn, T. Moscibroda, and R. Wattenhofer. What cannot be computed locally! In *23rd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 300–309. ACM Press, July 2004.
- [Lin87] N. Linial. Distributive graph algorithms - Global solutions from local data. In *28th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 331–335. IEEE Computer Society Press, October 1987.
- [Lin92] N. Linial. Locality in distributed graphs algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.
- [LWL⁺07] B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James, A. Gilchrist, and A. G. White. Experimental demonstration of a compiled version of shor’s algorithm with quantum entanglement. *Physical Review Letters*, 99(25):4, 2007.
- [Lyn97] N. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, 1997.
- [Mer90] N. D. Mermin. Quantum mysteries revisited. *American Journal of Physics*, 58(8):731–734, 1990.
- [Nao91] M. Naor. A lower bound on probabilistic algorithms for distributive ring coloring. *SIAM Journal on Discrete Mathematics*, 4(3):409–412, 1991.
- [NC00] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [Nie99] M. Nielsen. Conditions for a class of entanglement transformations. *Physical Review Letters*, 83(2):436–439, July 1999.
- [NS95] M. Naor and L. Stockmeyer. What can be computed locally. *SIAM Journal on Computing*, 24(6):1259–1277, 1995.

- [OMM04] M. Owari, K. Matsumoto, and M. Murao. Entanglement convertibility for infinite-dimensional pure bipartite states. *Physical Review A*, 70(5):1–4, 2004.
- [PCZ⁺08] J.-W. Pan, Z.-B. Chen, M. Żukowski, H. Weinfurter, and A. Zeilinger. Multi-photon entanglement and interferometry. Technical report, arXiv: quant-ph/0805.2853v1, May 2008.
- [Pel99] D. Peleg. Proximity-preserving labeling schemes and their applications. In P. Widmayer, G. Neyer, and S. Eidenbenz, editors, *25th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, volume 1665 of Lecture Notes in Computer Science, pages 30–41. Springer, June 1999.
- [PSK03] S. P. Pal, S. K. Singh, and S. Kumar. Multi-partite quantum entanglement versus randomization: Fair and unbiased leader election in networks. Technical report, arXiv: quant-ph/0306195v1, June 2003.
- [Sch05] M. Schlosshauer. Decoherence, the measurement problem, and interpretations of quantum mechanics. *Rev. Mod. Phys.*, 76(4):1267–1305, Feb 2005.
- [Shi84] A. Shimony. Controllable and uncontrollable nonlocality. In *International Symposium on the Foundations of Quantum Mechanics*, pages 130–139. The Physical Society of Japan, 1984.
- [Sho94] P. W. Shor. Algorithms for quantum computation: Discrete log and factoring. In *35th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 124–134. IEEE Computer Society Press, November 1994.
- [Sho97] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
- [SKP04] S. K. Singh, S. Kumar, and S. P. Pal. Characterizing the combinatorics of distributed EPR pairs for multi-partite entanglement. Technical report, arXiv: quant-ph/0306049v2, January 2004.
- [Str07] R. F. Streater. *Lost Causes in and Beyond Physics*. Springer-Verlag, 2007.
- [TKM05] S. Tani, H. Kobayashi, and K. Matsumoto. Exact quantum algorithms for the leader election problem. In *22nd Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 3404 of Lecture Notes in Computer Science, pages 581–592. Springer, February 2005.
- [Ż08] M. Żukowski. On Bell’s Theorem, quantum communication, and entanglement detection. In *Foundations of Probability and Physics 5*, August 2008.

A Model of a Quantum Distributed System

Any quantum-mechanical discussion relies on two fundamental concepts: *states* and *observables*. Intuitively, a state can be treated as a measure of knowledge about a physical system (usually associated with some observer), whereas the set of observables encodes the measurable properties of the system. There is a duality between these two concepts: we can also say that the state of the system is uniquely described through the distribution of outcomes of measurements on all possible observables related to it. Whereas in quantum-informational papers it is often convenient to focus on states, we choose to adopt the approach more usual in mathematical physics, which focuses on operator algebras (observables are operators satisfying certain mathematical conditions). The algebraic approach used as the basis of this model is generally accepted as the most mathematically robust theory, and moreover it naturally encodes concepts of quantum-mechanical locality, since operator algebras are spatially localised (restricted to each processor), unlike the state which is a global property of the system of all processors.

Most of the considerations in this paper can be without much loss of generality viewed as finite-dimensional. Then, a (*unital*) C^* -algebra can be introduced simply as some set of $m \times m$ matrices over complex numbers, which contains the identity matrix, and is closed with respect to the operations of matrix multiplication, matrix conjugation, and linear combination. Further on we rely only on a few basic concepts which can be understood in accordance with their standard definition for matrices, such as matrix multiplication, the tensor product \otimes , and spectral decomposition. The interested reader is referred to [BR79] for an explanation of more advanced concepts related to C^* -algebras.

A.1 Specification of the Physical System

The distributed set-up is given by a recipe which is well defined for each $n \in \mathbb{N}$. For simplicity of the description, we assume that all the processors know a value D , which is some arbitrarily weak upper bound on the number of neighbours of the node in the system, and that we consider problems for which the input and output values are integers also bounded by D , i.e., $x, y : V \rightarrow \{0, \dots, D\}$.

- Each processor v is described by its own copy \mathcal{A}_v of a C^* -algebra \mathcal{A} , localised in an area associated with the processor.
- The algebra \mathcal{A} describing a processor is a complex system composed of the processor's several modules, given in the form of the following tensor product of C^* -algebras: $\mathcal{A}_v = \mathcal{Q} \otimes \mathcal{IO}_1 \otimes \dots \otimes \mathcal{IO}_D$, where:
 - \mathcal{Q} is a non-commutative (quantum) algebra encoding the computational characteristics (“hardware”) of the processor,
 - $\mathcal{IO}_1, \dots, \mathcal{IO}_D$ are isomorphic copies of an algebra \mathcal{IO} representing a single input/output communication port of the processor (the algebra \mathcal{IO} is commutative if and only if the channel is classical, and non-commutative if and only if the channel is a quantum one).
- The multi-processor environment as a whole is described by the tensor product algebra of the algebras of specific processors, $\mathcal{A}^\otimes = \mathcal{A}_1 \otimes \dots \otimes \mathcal{A}_n$.

- The state of the multi-processor system is a positive normalized linear functional acting on algebra \mathcal{A}^\otimes , of the form $\omega : \mathcal{A}^\otimes \rightarrow \mathbb{R}_+$. The state ω is defined at the time of the initial set-up of the distributed system, and can be used to encode pre-entanglement.
- The *input data* of processors, given by way of a function $x : V \rightarrow \{0, \dots, D\}$, is fixed and for simplicity assumed to be *outside* the quantum system. The *algorithm* is defined by way of a family of quantum operations (completely positive maps), $\varphi^x : A \rightarrow A$, for $0 \leq x \leq D$, which encode the local operation of a processor having x as its input value.
- The evolution (dynamics) of the system is given through a sequence of discrete rounds. The t -th round is subdivided into a phase in which some local transformations φ_v are applied within each processor (computation phase), and a phase used for exchanging messages ψ_e along edges e between adjacent processors (communication phase). For an observable A^\otimes , initially we put $A_{(0)}^\otimes = \varphi_1 \cdots \varphi_n A^\otimes$, and for all subsequent rounds, $A_{(t+1)}^\otimes = \varphi_1 \cdots \varphi_n \psi_1 \cdots \psi_m A_{(t)}^\otimes$. The specific maps φ_v and ψ_e are formally defined as follows.
 - Dynamic maps $\varphi_v : \mathcal{A}^\otimes \rightarrow \mathcal{A}^\otimes$ describe local operations performed during a round t at each node. The map φ_v acts only on the local algebra of processor v depending on its input label $x(v)$, and is given by extension to the tensor product of the following transformation: $\varphi_v(A_1 \otimes \cdots \otimes A_v \otimes \cdots \otimes A_n) = A_1 \otimes \cdots \otimes \varphi^{x(v)}(A_v) \otimes \cdots \otimes A_n$. In each round t , all maps φ_v corresponding to different processors act on independent algebras and clearly commute (i.e., they can be executed simultaneously or reordered without changing the result).
 - Dynamic maps $\psi_e : \mathcal{A}^\otimes \rightarrow \mathcal{A}^\otimes$ describe communication along edges e of the system graph, and at the same time define the system graph. All the maps are induced by extension to the tensor product of the same transmission function for a pair of input/output ports $\psi : \mathcal{IO} \times \mathcal{IO} \rightarrow \mathcal{IO} \times \mathcal{IO}$, given simply as the exchange operation $\psi(X, Y) = (Y, X)$ if edge e exists in the graph, and the identity operation $\psi(X, Y) = (X, Y)$ otherwise (signifying lack of communication). The map ψ_e acts only on the the algebras corresponding to copies of \mathcal{IO} for the input/output ports of the processors communicating along edge e , leaving all other algebras unchanged. The maps ψ_e for different edges act on independent algebras and clearly commute (i.e., they can be executed simultaneously or reordered without changing the result).
- The local algebra \mathcal{A} of a processor contains one distinguished element: observable $M \in \mathcal{A}$, which is used for purposes of measurement. For simplicity we assume that M has a discrete spectral decomposition of the form $M = \sum_{i=1}^k \lambda^i P^i$, with eigenvalues $\lambda^i \in \mathbb{R}_+$ and projectors $P^i \in \mathcal{A}$. The observable $M^\otimes \in \mathcal{A}^\otimes$ is given through the tensor products of particular processors' copies of observable M , as $M^\otimes = M \otimes \cdots \otimes M$.
- The evolution of the system is assumed to terminate after T rounds. After T rounds, a standard von Neumann measurement process is applied to observable $M_{(T)}^\otimes = [M \otimes \cdots \otimes M]_{(T)}$ (the evolved observable M^\otimes after T rounds). This can be described as follows: by iterating over all possible tuples of values $(i_1, \dots, i_n) \in \{1, \dots, k\}^n$ for all v , with probability $p_{(i_1, \dots, i_n)} = \omega([P^{i_1} \otimes \cdots \otimes P^{i_n}]_{(T)})$, the values $(\lambda^{i_1}, \dots, \lambda^{i_n})$ are selected as the result of measurement for the respective processors $(1, \dots, n)$. The probabilities $p_{(i_1, \dots, i_n)}$ are understood here in the classical sense, and normalized in sum to 1.

- The output of the algorithm is obtained by applying a function $f : \mathbb{R}_+ \rightarrow \{0, \dots, D\}$ to the measurement results. Thus, values $(f(\lambda^{i_1}), \dots, f(\lambda^{i_n}))$ are returned by processors $(1, \dots, n)$, respectively.

Such a set-up is chosen for its simplicity, but obviously, there exist several other definitions which lead to equivalent models. For example, the input can be represented by enlarging the algebras \mathcal{Q} and introducing an additional factored part of the input state. Also, for improved clarity of the model, we prefer to consider dynamic maps, deferring other quantum operations (such as measurements) to the final stage of the algorithm. This can be achieved without loss of generality by sufficiently enlarging the local algebras of the processors (see e.g. [Str07] for a high-level exposition). When defining algorithms in practice it may of course be convenient to apply measurements in intermediate steps so as to simplify formulation.

In what follows, we introduce some standard notation. A state ω is called *pure* if it is extremal with respect to convex combination of states (i.e., if $\omega = \alpha\omega_1 + (1 - \alpha)\omega_2$ for some states $\omega_1 \neq \omega_2$ and $0 \leq \alpha < 1$, then $\alpha = 0$). A pure state ω is said to be a *product state* over $\mathcal{A}_a \otimes \mathcal{A}_b$ if for any $A \in \mathcal{A}_a, B \in \mathcal{A}_b$ we have $\omega(A \otimes B) = \omega(A \otimes \mathbf{1})\omega(\mathbf{1} \otimes B) \equiv \omega_a(A)\omega_b(B)$; for compactness, we simply write $\omega = \omega_a\omega_b$.

A.2 Details of the Computational Model

From a computational perspective, the set-up of the system described in the previous section can be summarized as follows:

- The *distributed algorithm* $(\omega, \{\varphi^x\}, \mathcal{Q}, \mathcal{IO}, M, f)$ is defined by setting the initial state ω of the system, the maps φ^x which shape the local computations in each round, and the observable M and function f responsible for extracting the output from the quantum system. The algebras \mathcal{Q} and \mathcal{IO} which define the “hardware” of the processors can be included in the specification of the algorithm, or can be taken as the general operator algebra $\mathcal{B}(H)$ over a Hilbert space.
- The *input* is provided by setting the edges of the system graph and the inputs $x(v)$ of specific nodes; these settings directly influence the maps ψ_e and φ_v which are responsible for communication and local computations within the system, respectively.

Definition 5. The above described model of a physical system provides a formal characterization of quantum extensions to the *LOCAL* model.

- When no restrictions are made about the state ω , the model is said to be equipped with the ${}^+\mathcal{E}$ extension.

When state ω is restricted to be a mixed state separable over the local algebras \mathcal{A} , i.e., a state of the form $\sum_i p_i(\omega_1^i \cdots \omega_n^i)$, for some values of probabilities $p_i, p_i \geq 0, \sum_i p_i = 1$, and some local pure states ω_n^i over algebra \mathcal{A} , possibly different for each processor, then the model is said to be equipped with the ${}^+\mathcal{S}$ extension.

When state ω is restricted to be of the form $\omega_l \cdots \omega_l$, for some pure state ω_l over algebra \mathcal{A} , identical for each processor, the system has neither of these extensions.

- When no restrictions are made about the communication algebra \mathcal{IO} , which may be non-commutative, the system is said to be equipped with the ${}^+\mathcal{Q}$ extension.

When algebra \mathcal{IO} is restricted to be commutative, communication is understood in the classical sense and the system has no such extension.

Note that the above definition also characterizes the \mathcal{LOCAL} and $\mathcal{LOCAL}^+\mathcal{S}$ models. The fact that such a characterisation is equivalent to the computational definition (Section 1.3) is straightforward to prove, taking into account that the considered initial states are separable, and the proposed evolution cannot create entanglement since only classical algebras are applied for interaction between processors (cf. [Bae87] and [Nie99] for different expositions of related concepts). The evolution on separable states is then easily simulated by a stochastic process, or equivalently, a distributed randomized algorithm.

B Proof of Theorem: $\mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E} \subseteq \varphi\text{-}\mathcal{LOCAL}$

Proof. Consider any quantum distributed algorithm $(\omega, \{\varphi^x\}, \mathcal{Q}, \mathcal{IO}, M, f)$. For any subset of processors $S \subseteq V$, we will denote by $\mathcal{A}^{\otimes S}$ the tensor product of algebras \mathcal{A} , taken over processors from set S , only. Consider the effect of the evolution maps $\{\varphi_v\}$ and $\{\psi_e\}$ on an arbitrary operator $A_S \in \mathcal{A}^{\otimes S}$. It is clear that A_S is not affected by φ_v if $v \notin S$, and moreover A_S is not affected by ψ_e if $e \notin \mathcal{V}_1(S)$. Thus, since these maps encode the input graph G_x , after a single round of evolution we may write $[A_S]_{(1)} = \tau_1(A_S, \mathcal{V}_1(G_x, S))$, where τ_1 denotes some operation dependent only on the algorithm (independent of G_x). Moreover, since maps φ_v are local and maps ψ_e only act on nearest neighbours, we have $[A_S]_{(1)} \in \mathcal{A}^{\otimes \mathcal{V}_1(S)}$. By applying the evolution procedure for t rounds, we immediately obtain by induction that $[A_S]_{(t)} = \tau_t(A_S, \mathcal{V}_t(G_x, S))$, where τ_t denotes some operation dependent only on the algorithm. Since the above holds for any operator $A_S \in \mathcal{A}^{\otimes S}$, distributions of results of all measurements restricted to set of processors S are independent of the input graph G_x , except for the local view $\mathcal{V}_t(G_x, S)$. This immediately implies that all solutions which belong to $\mathcal{LOCAL}^+\mathcal{Q}^+\mathcal{E}[t]$ also belong to $\varphi\text{-}\mathcal{LOCAL}[t]$. \square

C Supplementary Propositions

C.1 Properties of problem \mathcal{P}' (selecting an edge of $K_n \setminus \{e\}$)

Proposition 9. *Problem $\mathcal{P}' \notin \mathcal{LOCAL}[0, e^{-1}]$.*

Proof. Consider an outcome \mathcal{O} in $\mathcal{LOCAL}[0]$ which solves \mathcal{P}' with some probability Π . Within $\mathcal{LOCAL}[0]$, the output value $y_i \in \{0, 1\}$ of each node i is dependent only on the input label of the node, hence we may assume that with some probability p_i node i returns 0, and with probability $1 - p_i$ returns 1. From now on we will only consider the nodes which return 1 with non-zero probability, i.e. $p_i < 1$; w.l.o.g. let us suppose that this is the set of nodes $\{1, \dots, k\}$, for some k , $2 \leq k \leq n$. W.l.o.g. we can assume that $p_1 \leq p_2 \leq \dots \leq p_k$. It is straightforward to see that the worst-case input G_x for such a procedure is a graph with a missing edge between nodes 1 and 2. Hence, we can consider a simple application of Bernoulli's formula for a sequence of k independent random trials $(1, \dots, k)$ with failure probabilities (p_1, \dots, p_k) , in which the "winning event" is the success of exactly two trials, different from the pair $\{1, 2\}$. Denoting $q_i = p_i^{-1} - 1$, we have⁹:

$$\Pi = p_1 \cdots p_k \cdot \left(\sum_{1 \leq i < j \leq k} q_i q_j - q_1 q_2 \right). \quad (1)$$

It now suffices to prove that the above expression, treated as a multi-variable function with respect to k and $\{p_1, \dots, p_k\}$, subject to the constraints $k \in \{2, \dots, n\}$, $0 \leq p_1 \leq p_2 \leq \dots \leq$

⁹When $p_i = 0$, all subsequent expressions are well defined through their limit values when $p_i \rightarrow 0_+$.

$p_k < 1$, does not achieve a value $\Pi \geq e^{-1}$. This is a technical step which, for the sake of completeness, we perform below.

First, by directly solving a multi-variable optimization problem for small values of k , we establish that $\Pi < e^{-1} \approx 0.36$ for all $k \leq 5$. Indeed, when $k = 2$ we verify that $\Pi = 0$, when $k = 3$, $\Pi \leq \frac{8}{27} < 0.30$ (and $\Pi = \frac{8}{27}$ is attained when $p_1 = p_2 = p_3 = \frac{1}{3}$); when $k = 4$, $\Pi \leq \frac{5}{16} < 0.32$ (and $\Pi = \frac{5}{16}$ is attained when $p_1 = p_2 = p_3 = p_4 = \frac{1}{2}$), when $k = 5$, $\Pi \leq (\frac{3}{4})^4 < 0.32$ (and $\Pi = (\frac{3}{4})^4$ is attained when $p_1 = 0$ and $p_2 = p_3 = p_4 = p_5 = \frac{3}{4}$).

Now, let $k \geq 6$; we will show that it suffices to consider the case when $p_2 = p_3 = \dots = p_{k-1}$. Suppose, to the contrary, that $p_{k-1} > p_2$ and that $\Pi \geq e^{-1}$; we will construct an outcome \mathcal{O}' corresponding to some k' and some set of probabilities $\{p'_1, \dots, p'_k\}$, such that either $|\{i : p'_i = p'_2\}| > |\{i : p_i = p_2\}|$ or $k' < k$, and \mathcal{O}' yields a correct solution with probability $\Pi' \geq \Pi$. To achieve this, we put: $k' = k$, $p'_i = p_i$ for all $i < k - 1$, $p'_{k-1} = \max\{p_2, p_{k-1}p_k\} < p_{k-1}$, and $p'_k = \frac{p_{k-1}p_k}{p_{k-1}} > p_k$ (note that either $p'_{k-1} = p_2$, or else $p'_k = 1$; in the latter case, we set $k' = k - 1$). For convenience of notation we do not sort the values p'_i in non-decreasing order with respect to i , but note that we still have $p'_1 \leq p'_2 \leq p'_3, \dots, p'_k$, so Π' is given by an expression analogous to (1). We now introduce some auxiliary notation. For any subset of nodes $S \subseteq \{1, \dots, k\}$ with $|S| \geq 2$, let s_0, s_1 , and s_2 be the probabilities that exactly 0, 1, and 2 of the nodes from S , respectively, return a value of 1 in \mathcal{O} (s'_0, s'_1 , and s'_2 are likewise defined for \mathcal{O}' ; we also extend this notation to letters x for set X , and y for set Y). Denoting $\gamma_{(S,w)} = \sum_{i \in S} q_i^w$, for $w \in \{1, 2\}$, we recall the following simple relations:

$$\begin{aligned} s_0 &= \prod_{i \in S} p_i \\ s_1 &= \prod_{i \in S} p_i \cdot \sum_{i \in S} q_i = s_0 \gamma_{(S,1)} \\ s_2 &= \prod_{i \in S} p_i \cdot \sum_{i,j \in S, i < j} q_i q_j = \frac{1}{2} s_0 (\gamma_{(S,1)}^2 - \gamma_{(S,2)}) \leq \frac{s_1^2}{2s_0} \\ s_0 + s_1 + s_2 &\leq 1 \end{aligned} \tag{2}$$

Now, let $X = \{1, \dots, k - 2\}$ and $Y = \{k - 1, k\}$. By introducing the above relations into formula (1), we obtain for outcome \mathcal{O} :

$$\Pi = x_0 y_2 + x_1 y_1 + (x_2 - x_0 q_1 q_2) y_0. \tag{3}$$

When writing an analogous expression for \mathcal{O}' , we observe that $x'_0 = x_0$, $x'_1 = x_1$, $x'_2 = x_2$, $y'_0 = y_0$, $q'_1 = q_1$, and $q'_2 = q_2$, hence:

$$\Pi' = x_0 y'_2 + x_1 y'_1 + (x_2 - x_0 q_1 q_2) y_0.$$

By subtracting the above expressions we obtain and noting that $y_1 - y'_1 = -(y_2 - y'_2)$, we obtain:

$$\Pi' - \Pi = (x_1 - x_0)(y'_1 - y_1).$$

We have that $y'_1 - y_1 > 0$ (because $p_{k-1}p_k = p'_{k-1}p'_k$, $p'_{k-1} < p_{k-1}$, and $p'_k > p_k$). Consequently, if $x_1 \geq x_0$, then $\Pi' \geq \Pi$. We will now prove that the opposite case, i.e. $x_1 < x_0$ is impossible. Indeed, supposing that $x_1 < x_0$, consider the maximum possible value of y_0 , subject to the following constraints on probabilities $x_0, x_1, x_2, y_0, y_1, y_2$ which are then necessarily fulfilled: $y_0 + y_1 + y_2 = 1$, $x_0 + x_1 + x_2 \leq 1$, $x_1 \leq x_0$, $x_2 \leq \frac{1}{2}x_1$ (by (2)), $x_0 y_2 + x_1 y_1 + x_2 y_0 \geq e^{-1}$ (since the left-hand side is an upper bound on Π by (3), and moreover $\Pi \geq e^{-1}$ by assumption). By

solving this optimization problem, we obtain $y_0 \leq 1 - e^{-1}$ (with equality obtained for $x_0 = 1$ and $x_1 = x_2 = y_1 = 0$). Hence, since $y_0 = p_{k-1}p_k$, and $p_{k-1} \leq p_k$, we have $p_{k-1} \leq \sqrt{1 - e^{-1}} < \frac{4}{5}$. Consequently, for all $i \leq k - 2$, $p_i \leq p_{k-1} < \frac{4}{5}$, and $q_i = p_i^{-1} - 1 > \frac{1}{4}$. But then: $\frac{x_1}{x_0} = \gamma_{(X,1)} = q_1 + \dots + q_{k-2} > \frac{1}{4}(k-2) \geq 1$, a contradiction with the assumption $x_1 < x_0$.

Thus, it only remains to solve the case when $k \geq 6$ and $p_2 = p_3 = \dots p_{k-1}$. But then by developing expression (1), we have:

$$\Pi = p_2^{k-4}[(1-p_1)(1-p_k)p_2^2 + (k-3) \cdot (1-p_1)(1-p_2)p_2p_k + (k-2) \cdot (1-p_k)(1-p_2)p_1p_2 + \binom{k-2}{2}(1-p_2)^2p_1p_k].$$

This expression involves only four variables, subject to the constraints: $k \in \{6, 7, \dots, n\}$ and $0 \leq p_1 \leq p_2 \leq p_k \leq 1$, and by solving this problem we indeed verify that always $\Pi < e^{-1}$, which completes the proof. \square

We remark that the optimal assignment of probabilities to nodes is as follows: for $n = 3$ or $n = 4$, we put the uniform distribution, $p_i = \frac{n-2}{n}$. For $n \geq 5$, we always select 1 as the output for node 1 ($p_1 = 0$), and put $p_i = \frac{n-2}{n-1}$ for $2 \leq i \leq n$. As n tends to infinity, the probability that the obtained output is correct tends to e^{-1} from below. Clearly, in order to obtain such a solution, the nodes must take advantage of their input labels which contain information about the identifiers and n .

Proposition 10. *Problem $\mathcal{P}' \in \mathcal{LOCAL}^+\mathcal{S}[0, 1 - O(1/N^2)]$, when considering graphs of order $n \geq N$.*

Proof. It suffices to observe that the helper variable $h(v)$ can simply encode the output, and be set in a randomized way so as to select all edges of the complete graph with equal probability. For any input, the probability that the solution is correct is then $1 - 1/\binom{n}{2}$. \square

C.2 Characterization of the \mathcal{LOCAL} Model through Feasible Outcomes

For deterministic algorithms in \mathcal{LOCAL} , the output of any processor v after t rounds is a deterministic function of the identifiers of nodes which are located within its view $\mathcal{V}_t(G_x, v)$.

Claim 11. *An outcome $G_x \mapsto \{(y^i, p^i)\}$ can be deterministically obtained in t rounds in the \mathcal{LOCAL} model if and only if $p^1 = 1$ and there exists a function f such that, for every node v , $y^1(v) = f(\mathcal{V}_t(G_x, v))$.*

Allowing for randomised algorithms slightly increases the power of the model. The usual scenario is to allow each node v to flip coins during the execution of the algorithm, which is in fact equivalent to picking an arbitrary (possibly large) random positive integer $r(v)$ by each node before the start of execution [Lin92]. Outcomes in the \mathcal{LOCAL} model can thus be described by counting the number of random integer assignments leading to a particular output vector.

Claim 12. *An outcome $G_x \mapsto \{(y^i, p^i)\}$ belongs to $\mathcal{LOCAL}[t]$ if and only if there exists a function f such that for some integer k we have $p^i = \frac{1}{k^n} |\{r \in \{1, \dots, k\}^V : \forall v \in V, y^i(v) = f(\mathcal{V}_t(G_{(x,r)}, v))\}|$.*