

Distributed measurement-based quantum computation

Vincent Danos¹

Université Paris 7 & CNRS

Ellie D'Hondt²

Vrije Universiteit Brussel

Elham Kashefi³

IQC - University of Waterloo

Prakash Panangaden⁴

McGill University

Abstract

We develop a formal model for distributed measurement-based quantum computations, adopting an agent-based view, such that computations are described locally where possible. Because the network quantum state is in general entangled, we need to model it as a global structure, reminiscent of global memory in classical agent systems. Local quantum computations are described as measurement patterns. Since measurement-based quantum computation is inherently distributed, this allows us to extend naturally several concepts of the measurement calculus [2], a formal model for such computations. Our goal is to define an assembly language, i.e. we assume that computations are well-defined and we do not concern ourselves with verification techniques. The operational semantics for systems of agents is given by a probabilistic transition system, and we define operational equivalence in a way that it corresponds to the notion of bisimilarity. With this in place, we prove that teleportation is bisimilar to a direct quantum channel, and this also within the context of larger networks.

Key words: Formal language, quantum communication, quantum computing, semantics.

1 Introduction

Measurement-based models provide an intriguing new framework for thinking about quantum computation. While quantum circuits are still widely considered as a convenient formalism for describing algorithms, using measurements to steer quantum computation is considered a serious alternative. Due to their inherently probabilistic nature, measurements were long thought to be a disturbance to quantum computations – unavoidable though they are when wanting to read out the final output of a computation. That they can be an active component of a computation has been known for quite some time through the teleportation protocol. Only much later was it realized that in fault-tolerant constructions, measurements can be quite useful. Soon thereafter, with the advent of models such as the one-way quantum computer [11] and the teleportation model [9,8], it was established that measurements could not only be a recurring component of a computation, but the actual driving force behind it. Moreover, the measurement paradigm throws a whole new light on the strategies for actual physical implementations of a quantum computer [12].

However, measurements are not the only crucial ingredient of these models: they are also inherently *distributed*. Indeed, it is the realization that a variation on the teleportation protocol not only transports but also *transforms* quantum information, which is the basis of the teleportation model. Likewise, the one-way quantum computer is all about transformation via measurement and transportation, this time by way of a generic entangled state, the *graph state*. One-qubit measurements on this state transform the logical qubits, i.e. the quantum inputs, while transporting them via a path of graph state qubits. Again, non-local correlations provided by particular entanglement properties, together with measurements, steer the computation. Of course quantum measurements remain intrinsically probabilistic, but this can be solved by applying corrections dependent on previous measurement outcomes, rendering computations effectively deterministic. Note that measurement outcome dependencies are crucial in order to obtain universality of the model. As a result, the typically distributed notion of classical communication is also naturally present in these models. All of this is very nicely captured by the *measurement calculus* [2], a formal framework for one-way computations. Measurement patterns are defined essentially by sequences of commands allowed in the one-way model. From this one can define operational and denotational semantics and show their equivalence, and prove that notions of composition are well-defined. More importantly, there is an associated *rewrite system* which allows one to put any pattern into a standard form. The measurement calculus, which can

¹ Email: Vincent.Danos@pps.jussieu.fr

² Email: Ellie.DHondt@vub.ac.be

³ Email: ekashefi@iqc.ca

⁴ Email: Prakash@cs.mcgill.ca

be seen as an *assembly language*, proves to be a valuable tool for formal investigations into all measurement-based models; for example, one can easily show how the teleportation model reduces to the one-way model, via a conversion between the associated calculi [3].

Because of the inherently distributed aspect, measurement-based models for quantum computation are well-suited as a starting point for a formal model for distributed quantum computations. By this we mean *macroscopically* distributed, i.e. we are talking about coordinated actions between different parties. Of these there are many examples in quantum computation [10]: teleportation, of course, but also entanglement swapping, logic gate teleportation, cryptographic protocols, and also quantum versions of classical distributed applications such as leader election [5]. However, a formal language for distributed quantum computation is lacking. There have recently been interesting developments based on classical process calculi [7,6], which have focused mostly on the concurrency aspects. While the distributed nature of computations was introduced via types in Ref. [6], the aim there is to develop formal verification techniques. In this work, we take an assembly-language point of view, assuming that computations are well-defined. This results in a compact model, with which we can explore properties of distributed protocols, such as the coordination between agents.

We define an assembly language for distributed applications, directly built on the most basic distributed model of all: the one-way quantum computer. We adopt a *local* view and describe the system as a set of *agents* communicating synchronously and operating on a globally entangled quantum state; this is explained in Sec. 2. In Sec. 2.1 we develop a formal semantics for systems of agents as probabilistic transition systems. Operational equivalence is defined in a way such that it corresponds to the notion of bisimilarity. We then prove that quantum teleportation is bisimilar to a direct quantum channel, and this also within the context of other, possibly entangled, agents, in Sec. 3. While the correctness of teleportation has been proved within other formal frameworks before [6,1], the bisimilarity approach, and specifically, taking into account larger contexts as well, is new. We conclude in Sec. 4.

Some familiarity with the measurement calculus model is assumed; for an in-depth exposition we refer to Ref. [2].

2 Networks of agents

The main concept in our model for distributed measurement-based quantum computations is that of an *agent*. Agents are localized processes which, executing concurrently, make up a distributed system. Formally, we define agents in the following way.

Definition 2.1 An *agent* $\mathbf{A}(\mathbf{i}, \mathbf{o}) : Q, \mathcal{E}$, with classical input \mathbf{i} and output \mathbf{o} , and sort given by a set of qubit references Q , is defined by a finite event sequence \mathcal{E} composed of

- (i) patterns command sequences \mathcal{A} , with input qubits in Q ;
- (ii) classical message reception $c?x$ and sending $c!y$, where c is a classical channel, and x and y are names;
- (iii) qubit reception $qc?x$ and sending $qc!q$, where qc is a quantum channel and q a qubit reference.

An agent's *state* is given by a classical environment Γ , which is a partial mapping from names, i.e. classical variables and qubit references, to values.

Notice that any pattern $\mathcal{P}(V, I, O, \mathcal{A})$ trivially corresponds to agent $\mathbf{A} : I.\mathcal{A}$. In general the sort Q equals $I \uplus I^s$, where I is the local quantum input and I^s are qubits of a shared entangled state supplied by the network – see Def. 2.3 below. The classical input \mathbf{i} and output \mathbf{o} allows us to model protocols such as superdense coding, in which an agent wishes to send classical values to another agent with the help of a shared Bell state. The local state is used to store measurement outcomes resulting from local pattern executions, input from other agents over classical channels c , and classical input bindings from the local environment. Further bindings are added to the state as required; that is, for example, whenever a qubit is measured, only then the signal name s_i corresponding to the qubit reference q_i is added to the domain of Γ and bound to the classical measurement outcome v , denoted $\Gamma[s_i \mapsto v]$. The classical output set \mathbf{o} determines which bindings in Γ have to be preserved for the final output of a computation. We denote the local state restricted to the classical output by $\Gamma_{\mathbf{o}}$.

Our interpretation of agents is different than the usual process approach, in that agents in our setting always correspond to actual parties in a distributed network, denoted by the label \mathbf{A} . Therefore, an expression of the form $\mathbf{A}(\mathbf{i}, \mathbf{o}) : Q.\mathcal{E}$ should be read as: the agent with *name* \mathbf{A} (Alice) runs the *program* \mathcal{E} with qubits and in- and outputs as specified. An agents thus is a piece of code running on a particular processor. In this context, it only makes sense to compose agents if the agent names are the same. We refer to this as agent composition, which is formally defined below. Note that, since we allow subsequent agent programs to supply extra classical and quantum inputs as well as further process outputs from a previous program, some care is required in determining the classical type and quantum sort of the resulting agent. Other than that, agent composition is just straightforward concatenation of event sequences.

Definition 2.2 The composition of agents $\mathbf{A}(\mathbf{i}_1, \mathbf{o}_1) : Q_1.\mathcal{E}_1$ and $\mathbf{A}(\mathbf{i}_2, \mathbf{o}_2) : Q_2.\mathcal{E}_2$, which is only defined for agents with identical agent names, is denoted

$$\mathbf{A}[(\mathbf{i}_2, \mathbf{o}_2) : Q_2.\mathcal{E}_2] \circ [(\mathbf{i}_1, \mathbf{o}_1) : Q_1.\mathcal{E}_1]$$

and given by

$$\mathbf{A}(\mathbf{i}, \mathbf{o}) : Q.\mathcal{E}_2\mathcal{E}_1 \quad \text{with} \quad \begin{cases} \mathbf{i} = \mathbf{i}_1 \cup (\mathbf{i}_2 \setminus \mathbf{o}_1) \\ \mathbf{o} = \mathbf{o}_1 \cup \mathbf{o}_2 \\ Q = Q_1 \cup (Q_2 \setminus Q'_1), \end{cases} \quad (1)$$

where Q'_1 is the output sort of the first agent.

Note that output sorts can be determined by inspecting an agent's event sequence. The general idea is that $\mathbf{o}_1 \subseteq \mathbf{i}_2$ and $Q'_1 \subseteq Q_2$, so that subsequent programs are defined at least on the outputs of previous ones. This can always be arranged by assuming identity transformations for those names in \mathbf{o}_1 and Q'_1 not appearing in \mathcal{E}_2 . In particular this means that $Q' = Q'_2$. Extra inputs in subsequent programs are encountered for example when several qubits are teleported one after each other. In this case, Alice needs to execute her side of the protocol several times in a row, supplying a new local quantum input every time she initiates the protocol. The same argument holds for classical input when sequencing several dense coding protocols. Of course this is best understood within the context of *network* rather than agent composition, which is defined below.

A network of agents consist of several agents executing their event sequence concurrently, together with a global shared entangled state. As the network quantum state is inherently non-local, there is no other option than to regard it as some kind of global memory – even though we wish to adhere to a local view. This leads to the following definition.

Definition 2.3 A *network of agents* \mathcal{N} is defined by a set of concurrently acting agents together with a shared quantum state, that is

$$\begin{aligned} \mathcal{N} &= \mathbf{A}_1(\mathbf{i}_1, \mathbf{o}_1) : Q_1.\mathcal{E}_1 \mid \dots \mid \mathbf{A}_m(\mathbf{i}_m, \mathbf{o}_m) : Q_m.\mathcal{E}_m \parallel \sigma \\ &= \mid_i \mathbf{A}_i(\mathbf{i}_i, \mathbf{o}_i) : Q_i.\mathcal{E}_i \parallel \sigma, \end{aligned} \tag{2}$$

where $\sigma \in \mathcal{D}(\mathcal{H}_{\cup_i I_i^s})$, with $Q_i = I_i \cup I_i^s$ for all i .

The network state σ in the definition is the initial entanglement resource which is distributed among agents. Local quantum inputs specified in I_i are added to the network state σ during initialization. In this way we can keep initial shared entanglement as a first-class primitive in our model. In this paper, we do not describe the actual procedure for producing σ . Note that agents in a network need to have different names, since they correspond to different parties that make up the distributed system. In other words, concurrency comes *only* from distribution; we do not consider parallel composition of processes in the context of one party. Finally, individual agents $\mathbf{A}(\mathbf{i}, \mathbf{o}) : Q.\mathcal{E}$ trivially correspond to a network $\mathbf{A}(\mathbf{i}, \mathbf{o}) : Q.\mathcal{E} \parallel \mathbf{0}$. Therefore statements about networks affect individual agents as well.

We define two different ways of composing networks of agents, namely sequential and parallel composition. Because of our interpretation of agents as distributed processes, there are some constraints on these operations. Sequential composition is only defined for networks containing the same agents; the idea is that agents carry out event sequences of both networks one after the other. Furthermore, agent composition must be defined as per Def. 2.2, that is, inputs and initial sorts of the second network must contain outputs and

final sorts of the first. Formally, as follows.

Definition 2.4 The sequential composition of networks $\mathcal{N}_1 = \big|_{i=1}^m \mathbf{A}_i(\mathbf{i}_{1,i}, \mathbf{o}_{1,i}) : Q_{1,i} \cdot \mathcal{E}_{1,i} \parallel \sigma_1$ and $\mathcal{N}_2 = \big|_{i=1}^m \mathbf{A}_i(\mathbf{i}_{2,i}, \mathbf{o}_{2,i}) : Q_{2,i} \cdot \mathcal{E}_{2,i} \parallel \sigma_2$ is defined as

$$\mathcal{N}_2 \circ \mathcal{N}_1 = \big|_{i=1}^m \mathbf{A}_i[(\mathbf{i}_{2,i}, \mathbf{o}_{2,i}) : Q_{2,i} \cdot \mathcal{E}_{2,i}] \circ [(\mathbf{i}_{1,i}, \mathbf{o}_{1,i}) : Q_{1,i} \cdot \mathcal{E}_{1,i}] \parallel \sigma_1 \otimes \sigma_2. \quad (3)$$

Note that we have overloaded the notation \circ to denote both agent and network composition. As long as both networks have the same number of agents, one can always arrange for them two be sequentially composable by renaming agents. In fact one can even compose networks with a different number of agents by adding null agents, i.e. just agent names, to the network with fewer agents. The preparations of both networks are assumed to be defined on disjoint Hilbert spaces; the composed network's preparation is given by the tensor product of these. They are to be interpreted as different entanglement resources that are used by both networks. Parallel composition, which expresses that two networks are operating in parallel and independent of each other, is only defined for networks containing different agents. Again, one can always rename agents so that this is well-defined.

Definition 2.5 The composition of networks $\mathcal{N}_1 = \big|_{i=1}^m \mathbf{A}_i(\mathbf{i}_{1,i}, \mathbf{o}_{1,i}) : Q_{1,i} \cdot \mathcal{E}_{1,i} \parallel \sigma_1$ and $\mathcal{N}_2 = \big|_{i=1}^n \mathbf{B}_i(\mathbf{i}_{2,i}, \mathbf{o}_{2,i}) : Q_{2,i} \cdot \mathcal{E}_{2,i} \parallel \sigma_2$ is defined as

$$\begin{aligned} \mathcal{N}_1 \otimes \mathcal{N}_2 = & \big|_{i=1}^m \mathbf{A}_i(\mathbf{i}_{1,i}, \mathbf{o}_{1,i}) : Q_{1,i} \cdot \mathcal{E}_{1,i} \\ & \big|_{i=1}^n \mathbf{B}_i(\mathbf{i}_{2,i}, \mathbf{o}_{2,i}) : Q_{2,i} \cdot \mathcal{E}_{2,i} \\ & \parallel \sigma_1 \otimes \sigma_2. \end{aligned} \quad (4)$$

By combining sequential and parallel composition, one can express a broad range of network combinations. In Sec. 2.3 we show that the semantics of networks is preserved under these operations.

The definitions given in the above can be rendered concisely under the form of an abstract grammar. We use \square instead of $|$ to separate choices for expressions, as the latter is already in use to denote parallel composition.

$$\begin{aligned} \mathcal{A} & ::= \text{nil} \square E \square M \square C \square \mathcal{A} \otimes \mathcal{A} \square \mathcal{A} \cdot \mathcal{A} \\ \mathcal{E} & ::= \mathbf{c}?x \square \mathbf{c}!x \square \mathbf{qc}?q \square \mathbf{qc}!q \square \mathcal{A} \square \mathcal{E} \cdot \mathcal{E} \\ \mathbf{a} & ::= \mathbf{A}(\mathbf{i}, \mathbf{o}) : Q \cdot \mathcal{E} \square \mathbf{A}[(\mathbf{i}_2, \mathbf{o}_2) : Q_2 \cdot \mathcal{E}_2] \circ [(\mathbf{i}_1, \mathbf{o}_1) : Q_1 \cdot \mathcal{E}_1] \\ \mathcal{N} & ::= \big|_{i \mathbf{a}_i} \parallel \sigma \square \mathcal{N} \circ \mathcal{N} \square \mathcal{N} \otimes \mathcal{N} \end{aligned} \quad (5)$$

In the above E , M and C stand for any entanglement, measurement or correction commands, and nil is the null command. We refrain from giving a grammar for names; rather, we use conventions for these described below. Notice that both pattern command sequences and networks can be viewed as processes in the traditional process algebra sense, with a composition operation transforming any two patterns, respectively networks, into a new pattern

or network. The definition of event sequences, which can be composed sequentially, is also common. Agents, however, have a less clear status in the process algebra framework. This is exactly because they are constructs that formalize *distributed* notions, and therefore there are constraints on how they may be composed. This is not a flaw of our model but rather a requirement if one wants to make distribution explicit.

There are several notational conventions we adhere to throughout this paper. When the context is clear, as is often the case, we do not explicitly mention inputs and outputs, but instead just write $\mathbf{A} : Q$. Moreover, we do not write inputs, outputs, sorts, or preparations if there are none, sometimes writing $-$ for empty input or output sets. For example, $\mathbf{A}(-, \{x\}).c?x$ is an agent with no input and no qubits, while $\mathbf{A}(-, \{x\}).c?x \mid \mathbf{B}(\{y\}, -).c!y$ is a network with no preparation in which a classical value is exchanged between two agents. We write $c?xy$ for $c?xc?y$ and similarly for other communicating channels. In Def. 2.1 we have used pattern command sequences rather than full patterns for brevity, with the convention that the input of a pattern is always given by those qubits in its computation space that belong to the agent's sort at the moment when the pattern is executed. The output of the pattern is simply given by those qubits that are not measured. For example, in the single-agent network $\mathbf{A} : \{1, 2\}.X_4^{s_1}M_1^0E_{14}$ qubit 1 is an input to the pattern and 4 is an output; we also write this as $\mathbf{A} : \{1, 2\}.\mathcal{H}(1, 4)$. Pattern qubits not in an agent's sort are assumed to be computation qubits initialized to $|+\rangle$ as before, and need not be mentioned explicitly in an agent's sort. On the other hand, agent qubits not mentioned in a pattern event, as qubit 2 above, are always assumed to be left alone. That is, we do not explicitly write the identity pattern \mathcal{I} applied to \mathbf{A} 's remaining qubit 2. We consider pattern command *sequences* rather than singular commands so that we can use the big-step semantics of patterns. As can be seen from these examples, we often refer to qubits via numbers; specifically we mostly refer to qubit q_i as qubit i , and s_i instead of s_{q_i} for the corresponding signal variable. This has the advantage that patterns look just as they do in Ref. [2]. In what follows below, we use specific letters for specific names, in particular we use q_i or just i for qubit references, x_i and y_i for ordinary classical variables, v_i for classical values, and s_i for classical signal variables.

We subject networks of agents to definiteness conditions, which ensure that the computation is well-defined.

- (H0) an agent's communication events operate only on qubits in its sort Q .
Pattern events with computation space V have inputs in $Q \cap V$;
- (H1) an agent's events depend only on values in its state Γ ;
- (H2) each quantum and classical message reception event has a corresponding quantum, respectively classical message sending event;
- (H3) all names, i.e. classical variables and quantum references, are unique.

2.1 Operational semantics

Before we provide concrete evaluation rules for distributed computations, we give some clarifying explanations and examples as to how execution proceeds in the local view. Throughout network evolution, each agent has access to the network state σ via the qubits it owns, transforming σ whenever a pattern event is executed. While these patterns are local, σ is *not*, and to preserve all information on the correlations we need to keep σ unreduced at all times. As an example, suppose that an agent \mathbf{A} owns the first two qubits of the system's state $\sigma(1, 2, 3, 4, 5)$. Its next event is to execute the Hadamard pattern $\mathcal{H}(1, 6)$ on its first qubit. As stated above, we do not explicitly write the identity pattern \mathcal{I} applied to \mathbf{A} 's remaining qubit 2. More importantly, neither do we write explicit identity patterns for the qubits not belonging to \mathbf{A} , in this case 3, 4 and 5. In full, this means that we have an evaluation step as follows

$$\sigma, \mathbf{A} : \{1, 2\}. \mathcal{H}(1, 6) \Longrightarrow \sigma', \mathbf{A} : \{6, 2\}, \quad (6)$$

where $\sigma' = \llbracket \mathcal{H}(1, 6) \otimes \mathcal{I}^{\otimes 4}(2, 3, 4, 5) \rrbracket \sigma$, that is, $\sigma' = (H \otimes I^{\otimes 4})\sigma(H \otimes I^{\otimes 4})$. We denote the transition relation by \Longrightarrow . Execution of \mathcal{H} occurs in a single transition step by relying on its big-step semantics. In this way we avoid getting into the actual details of pattern execution, which is not what this paper is about. Note especially that the sort of \mathbf{A} has changed. This is because measurements are destructive, so that the input qubit 1 has disappeared, and the output 6 has taken its place. Because of this, we sometimes explicitly write $\sigma'(6, 2, 3, 4, 5)$ in the right-hand side of the above evaluation rule. The only remnant of 1 is its corresponding measurement outcome s_1 , which is recorded in the state Γ , via the added binding $\Gamma[s_1 \mapsto v]$, where v is the measurement outcome. \mathbf{A} might subsequently send s_1 to other agents by an event $\mathbf{c}!s_1$, so in general we cannot delete this entry from Γ . Essentially this means that when new qubit references are generated by the execution of subsequent local patterns, or by quantum communications from other agents, these names need to be *unique*. Because we do not want to get into the actual details of this naming procedure, as we are considering our model to be at the level of an assembly language, i.e. at a stage where naming conflicts have been resolved, we have imposed this as a definiteness condition in the above.

A final point concerns an agent's classical input and output, received from, respectively sent to, its own local system. In the local view, pattern events can depend also on classical inputs, rather than only on measurement outcomes. Because of the uniform structure of the local state Γ we can piggyback input dependencies onto the signal dependency structure of MC. This is, as mentioned before, one of the reasons why MC is such a good basis on which to build a framework for distributed quantum computations. With these concrete examples in mind, we are now ready to develop the operational semantics of the local view. As usual, we first define rules in terms of small-step transitions, after which we switch over to the big-step framework.

The small-step transitions for distributed computations essentially describe how agents, and the network with them, evolve over different time steps. We adopt a shorthand notation for agents, leaving out classical inputs and output, which do not change with small-step reductions.

$$\begin{aligned}
 \mathbf{a}_i &= \mathbf{A}_i : Q_i \cdot \mathcal{E}_i \\
 \mathbf{a}_i \cdot E &= \mathbf{A}_i : Q_i \cdot [\mathcal{E}_i \cdot E] \\
 \mathbf{a}^{-q} &= \mathbf{A} : Q \setminus \{q\} \cdot \mathcal{E} \\
 \mathbf{a}^{+q} &= \mathbf{A} : Q \uplus \{q\} \cdot \mathcal{E}[q/x],
 \end{aligned} \tag{7}$$

where E is some event, and \mathcal{E}_i and \mathcal{E}'_i are event sequences. A *configuration* is given by the system state σ together with a set of agent programs, and their states, specifically

$$\sigma, |_i \Gamma_i, \mathbf{a}_i = \sigma, \Gamma_1, \mathbf{a}_1 \mid \Gamma_2, \mathbf{a}_2 \mid \dots \mid \Gamma_m, \mathbf{a}_m. \tag{8}$$

The small-step rules for configuration transitions, denoted \Longrightarrow , are specified below; we give some explanations afterwards. When the system state is not changed in an evaluation step, we stress this by preceding a rule by $\sigma \vdash$.

$$\frac{\sigma, \mathcal{P}(V, I, O, \mathcal{A}) \longrightarrow_\lambda \sigma', \Gamma'}{\sigma, \Gamma, \mathbf{A} : I \uplus R \cdot [\mathcal{E} \cdot \mathcal{P}] \Longrightarrow_\lambda \sigma', \Gamma \cup \Gamma', \mathbf{A} : O \uplus R \cdot \mathcal{E}} \tag{9}$$

$$\frac{\Gamma_2(y) = v}{\sigma \vdash (\Gamma_1, \mathbf{a}_1 \cdot c?x \mid \Gamma_2, \mathbf{a}_2 \cdot c!y \Longrightarrow \Gamma_1[x \mapsto v], \mathbf{a}_1 \mid \Gamma_2, \mathbf{a}_2)} \tag{10}$$

$$\frac{}{\sigma \vdash (\Gamma_1, \mathbf{a}_1 \cdot qc?x \mid \Gamma_2, \mathbf{a}_2 \cdot qc!q \Longrightarrow \Gamma_1, \mathbf{a}_1^{+q} \mid \Gamma_2, \mathbf{a}_2^{-q})} \tag{11}$$

$$\frac{L \Longrightarrow_\lambda R}{L \mid L' \Longrightarrow_\lambda R \mid L'} \tag{12}$$

Here, \cup denotes the union of outcome maps. Implicit in these rules is a sequential composition rule, which ensures that all events in an agent's event sequence are executed one after the other. The first rule is for local operations; we have written the full pattern instead of only its command sequence here to make pattern input and output explicit. Because a pattern's big-step semantics is given by a probabilistic transition system described by \longrightarrow , we pick up a probability λ here. Furthermore, an agent changes its sort depending on pattern's output O , as explained in the examples above. The next rule is for classical rendez-vous and is straightforward. For quantum rendez-vous, we need to substitute q for x in the event sequence of the receiving agent, and furthermore adapt qubit sorts. The last rule is a meta-rule, which is required to express that any of the other rules may fire in the context of a larger system. L and R stand for any of the possible left-, respectively right-hand sides of any of the previous rules, while L' is an arbitrary configuration. Note that we

might need to rearrange terms in the parallel composition of agents in order to be able to apply the context rule. This can always be done since the order of agents in a configuration is arbitrary. In derivations of network execution, we often do not explicitly write reductions as specified by (12), but rather specify which in which order the other rules fire in the context of the network at hand. It is precisely in this last rule that introduces non-determinism at the network level, that is, several agent transitions may be possible within the context of a network at the same time.

Starting from the small-step rules above we can now define the big-step semantics of a system of agents. We first define computation paths, which run from initial to final configurations via small-step transitions. In the *initial configuration*, all local states are given by the map containing the classical input bindings $\Gamma_{\mathbf{i}_i}$, while the network state is determined by the entanglement resource together with local quantum inputs. A *final configuration* is one in which all agents have an empty command sequence, and in which all local states have been restricted to classical output bindings, $\Gamma_{\mathbf{i}, \mathbf{o}_i}$. Because of the definiteness conditions we imposed each computation always ends with a final configuration. Supposing the initial quantum state of agent \mathbf{A}_i is given by $\rho_i \in I_i$, paths are defined as follows.

Definition 2.6 Given a network of agents $\mathcal{N} = |_i \mathbf{A}_i(\mathbf{i}_i, \mathbf{o}_i) : Q_i \cdot \mathcal{E}_i || \sigma$ and quantum inputs ρ_i , a *path* γ is a *maximal* sequence of configurations $\{C_j = \sigma_j, |_i \Gamma_{\mathbf{i}_i}^j, \mathbf{a}_i^j, j = 1, \dots, k-1\}$, i.e.

$$\begin{aligned} C_1 &= \sigma \otimes_{i=1}^m \rho_i, |_i \Gamma_{\mathbf{i}_i}, \mathbf{A}_i : Q_i \cdot \mathcal{E}_i \\ C_j &\Longrightarrow_{\lambda_j} C_{j+1} \\ C_k &= \sigma_k, |_i \Gamma_{\mathbf{i}, \mathbf{o}_i}^k, \mathbf{A}_i : Q_i^k \end{aligned} \tag{13}$$

We write $C_1 \xrightarrow{\gamma} C_k$ where $\lambda_\gamma = \prod_{j=1}^k \lambda_j$, and call C_k a final configuration of \mathcal{N} .

Notice also that paths always terminate since event sequences are finite.

One could straightforwardly define the operational semantics of a system \mathcal{N} to be the probabilistic transition system (PTS) defined by of all its paths. However, we choose to identify those paths leading to the same observable behavior of an agent network. Concretely, for particular inputs \mathbf{i} and $\rho \in I$, we identify final configurations for which only internal bindings in the local state of agents are different, that is, bindings for names not in the classical output set \mathbf{o} . These bindings correspond either to outcomes of measurements appearing in pattern events, or result from classical rendez-vous events. As long as these are not part of the classical output, their actual values are unimportant. We cannot trace out these measurement outcomes after each pattern event, since a subsequent event may depend on these values. Notice that the final sorts of agents *do* need to be identical, as well as, obviously, the final network quantum state. Identifying such paths, in the style of the measure-

ment calculus, then gives us the semantics of a network of agents. However, because of the non-determinism in the order in which concurrent agents execute their event sequence, we have to define the operational semantics of a network with respect to a particular *schedule*. A schedule is precisely a particular order in which agents execute events. For example, in the network $\mathbf{A} : \{1\}.\mathcal{H}(1, 2) \mid \mathbf{B} : \{3\}.\mathcal{H}(3, 4)$, possible schedules are \mathbf{AB} and \mathbf{BA} . If we do not take schedules into account, we would add probabilities of all paths and all schedules resulting in identical final configurations, which for the example above leads to a probability of 2 for computing $H \otimes H$ for arbitrary inputs, which is clearly not what we intend to say. We refrain from giving a formal definition for schedules, as we will find in Sec. 2.2 that the semantics of a network is independent of the schedule. Putting all this together, we obtain the following definition.

Definition 2.7 The *operational semantics* of a network $\mathcal{N} = |_i \mathbf{A}_i(\mathbf{i}_i, \mathbf{o}_i) : Q_i.\mathcal{E}_i \parallel \sigma$, with respect to a particular schedule, is a probabilistic transition system relating initial with final configurations,

$$\llbracket \mathcal{N} \rrbracket_{op} : \uplus_i Q_i \rightarrow \uplus_i Q'_i \cdot \otimes_i \rho_i, |_i \Gamma_{\mathbf{i}_i} \Longrightarrow_{\lambda} \sigma', |_i \Gamma_{\mathbf{o}_i} \quad (14)$$

with $\lambda = \sum_{\gamma} \lambda_{\gamma}$ and the sum runs over all paths γ such that

$$\sigma \otimes_i \rho_i, |_i(\Gamma_{\mathbf{i}_i}, \mathbf{A}_i(\mathbf{i}_i, \mathbf{o}_i) : Q_i.\mathcal{E}_i) \xrightarrow{\gamma}_{\lambda_{\gamma}} \sigma', |_i(\Gamma_{\mathbf{o}_i}, \mathbf{A}_i(\mathbf{i}_i, \mathbf{o}_i) : Q'_i). \quad (15)$$

We call $\uplus_i Q_i \rightarrow \uplus_i Q'_i$ the *type* of the network.

From now on, we denote $I = \uplus_i I_i$ for the set of input qubits and $O = \uplus_i O_i$ for the set of output qubits, and call $\mathcal{D}(\mathcal{H}_I)$ and $\mathcal{D}(\mathcal{H}_O)$ the quantum input and output space respectively. The semantics of a network with respect to a schedule is thus that it relates quantum states in $\mathcal{D}(\mathcal{H}_I)$ plus classical input to quantum states in $\mathcal{D}(\mathcal{H}_O)$ and classical output with particular probabilities. Note that the type of the transition system is a mapping from initial to final sorts; this component is identical in the denotational semantics we develop in Sec. 2.2.

We say that two networks \mathcal{N}_1 and \mathcal{N}_2 are *operationally equivalent* if their operational semantics, given by a PTS, is identical, and write $\llbracket \mathcal{N}_1 \rrbracket_{op} \equiv \llbracket \mathcal{N}_2 \rrbracket_{op}$. In fact, we identify operational equivalence with the notion of *bisimilarity*. This is indeed sensible since, by identifying computation paths as in Def. 2.7, we actually impose a bisimilarity relation on final configurations. As we shall see in Sec. 3, it is by doing exactly this that we can show, among others, that teleportation is bisimilar to a direct quantum channel.

2.2 Denotational semantics

Because of its more abstract character, in many situations it is more adequate to work with the denotational semantics. This is why we develop this notion

for networks of agents. As before, it is closely related to the operational semantics, as well as to the semantics of ordinary patterns. Indeed, upon inspection of Def. 2.7, we see that for any schedule, the PTS associated with a network of agents decomposes in several parts. First, there is a map from initial to final sorts, which determines how qubit ownership evolves for each of the agents from initial to final configurations. This is formalized as a type signature, exactly as we did for the operational semantics. The sort mapping is independent of the classical input: indeed, classical inputs appear only in classical communications, Pauli corrections and measurement angles, none of which affect qubit sorts. Furthermore, they can be read off statically from the network definition and are also schedule independent. The denotational semantics is then a mapping from classical inputs to classical outputs and a quantum operation, which in turn determines how quantum states evolve in the network. However, these two components are not independent of each other, since classical outputs can be measurement outcomes, which occur with probabilities that depend on the quantum operation applied. For simplicity, let us first consider the case where there are no classical outputs. In this case, for each classical input $\mathbf{i} = \uplus_i \mathbf{i}_i$ we have map \mathcal{L} which describes how the network quantum state evolves. This map is a multilocal quantum operation, because if we throw away all distributed information, that is, sorts and communication events, we just have an ordinary pattern, i.e. a quantum operation. There is one caveat: since computation occurs asynchronously, there is usually some choice in the order in which different agents execute events in their program, i.e. there are different possible schedules. However, since at each instance of the computation local events operate on disjoint sets of qubits, it does not actually matter in which order these operations are applied, or, in fact, whether they are executed at the same time. This statement is proved formally in Sec. 2.4; we postpone the full proof until then since it has bearing on other situations that are covered below. Therefore, any schedule of the computation leads to the same quantum operation. So to determine the operation elements of \mathcal{L} , we choose a particular schedule, and then compose patterns in the order in which they are executed, tensoring with identity patterns where necessary and ignoring communication commands. Each operation element L_j then corresponds to a sequential composition of actualizations for each of these patterns.

Suppose now that the network contains classical outputs $\mathbf{o} = \uplus_i \mathbf{o}_i$. We need to make a distinction between *signal* outputs, which are measurement outcomes, and *external* outputs, which are values that were originally input by some agent and sent around the network. By definition, the external outputs $\mathbf{o}_e = \uplus_i \mathbf{o}_{i,e}$ depend only on the classical input \mathbf{i} ; these constant values are sent around the network via classical channels. It is precisely the signal outputs $\mathbf{o}_s = \uplus_i \mathbf{o}_{i,s}$ that depend on the quantum operation and vice versa. Indeed, when there are signal outputs particular measurement outcomes are preserved, therefore excluding actualizations of \mathcal{L} that do not correspond to

that outcome. This essentially means that a different quantum operation is applied for each possible signal output. For example, suppose one of the signal outputs, corresponding to a measurement on qubit 3 is equal to 1. Then only those actualizations containing the operator $\langle -_\alpha |_3$ are compatible with this output. We denote actualizations compatible with output \mathbf{o}_s by $L_i^{\mathbf{o}_s}$, and the quantum operation with these operation elements by $\mathcal{L}^{\mathbf{o}_s}$, and call these *restricted*. Note, however, that this is a *trace-decreasing* operation, and that $\text{Tr}(\mathcal{L})^{\mathbf{o}_s(\rho)}$ is precisely the probability with which the output \mathbf{o}_s occurs. So, whereas the operational semantics gives explicit probabilities for each path, in the denotational semantics these are contained within the quantum operations. It is this abstraction, together with schedule-independence which makes the denotational framework advantageous. Indeed, classical inputs are the same for all schedules, and classical outputs, depending on classical inputs and measurement values, thus occur with the same probabilities for all schedules, since \mathcal{L} is schedule-independent. Putting all of this together, we arrive at the following definition.

Definition 2.8 The *denotational semantics* of a network of agents $\mathcal{N} = |_i \mathbf{A}_i(\mathbf{i}_i, \mathbf{o}_i) : Q_i.\mathcal{E}_i || \sigma$ is given by

$$\llbracket \mathcal{N} \rrbracket_{de} : \uplus_i Q_i \rightarrow \uplus_i Q'_i.\mathbf{i} \rightarrow \{(\mathbf{o}, \mathcal{L}^{\mathbf{o}_s}), \forall \mathbf{o}_s\} \quad (16)$$

with

$$\mathcal{L} : \mathcal{D}(\mathcal{H}_I) \rightarrow \mathcal{D}(\mathcal{H}_O) : \otimes_i \rho_i \rightarrow \sum_j L_j(\sigma \otimes_i \rho_i) L_j^\dagger, \quad (17)$$

where $\mathbf{o} = \mathbf{o}_e \uplus \mathbf{o}_s$, ρ_i is the quantum input, Q'_i the final sort of agent \mathbf{A}_i , and I and O are quantum input and output spaces respectively. In case there are no outputs, we have $\llbracket \mathcal{N} \rrbracket_{de} : \uplus_i Q_i \rightarrow \uplus_i Q'_i.\mathbf{i} \rightarrow \mathcal{L}$, or just \mathcal{L} if there are no inputs either.

Note that the \mathbf{o}_e part of \mathbf{o} in each of the above tuples is identical. As an example, consider the pattern $X_1^{s_2} M_2^{-\alpha}$, which implements the bit-flip channel [10]. It can be interpreted as a one-agent network $\mathbf{A}(-, -) : \{1\}.X_1^{s_2} M_2^{-\alpha}$, with has as denotational semantics the quantum operation $\mathcal{L}(\rho) = p\rho + (1 - p)X\rho X$. However, the one-agent network $\mathbf{A}(-, \{s_2\}) : \{1\}.X_1^{s_2} M_2^{-\alpha}$ has a different semantics, namely

$$\{(0, p\rho), (1, (1 - p)X\rho X)\}, \quad (18)$$

for all ρ , where p is a function of α . While this example may seem contrived, it is actually crucial that the semantics of these kind of networks are different, as they describe different states of *knowledge* of the output s_2 , and hence, of what actual computation path was taken.

In case the underlying quantum operation \mathcal{L} is deterministic, all actualizations lead to the same quantum output. However, even in that case we require the above formulation with different trace-decreasing quantum operations $\mathcal{L}^{\mathbf{o}_s}$, since these determine the probabilities of outputting \mathbf{o}_s . A network

is deterministic only if for any input \mathbf{i} , the quantum operation \mathcal{L} is deterministic, i.e. it implements a unitary, *and* the classical output \mathbf{o} is identical in all actualizations.

As before, two networks are called denotationally equivalent if they have the same denotational semantics. Note that equivalent networks may have different preparation states, and also that agent names may be different, though the number of agents must be identical in both networks. With the above definitions in place, we can now prove the following result.

Proposition 2.9 *There is a precise correspondence between the operational and the denotational semantics of networks of agents, that is to say*

$$\forall \mathcal{N}_1, \mathcal{N}_2 : \mathcal{N}_1 \equiv_{op} \mathcal{N}_2 \iff \mathcal{N}_1 \equiv_{de} \mathcal{N}_2 \quad (19)$$

Proof. Suppose $\llbracket \mathcal{N}_1 \rrbracket_{de} = \llbracket \mathcal{N}_2 \rrbracket_{de}$. This means that for all classical and quantum inputs, both the the type and classical external outputs are identical. Since for all signal outputs $\mathcal{L}_1^{\mathbf{o}_s} = \mathcal{L}_2^{\mathbf{o}_s}$, we have that $\mathcal{L}_1 = \mathcal{L}_2$ and therefore computation paths are also the same for both networks [4]. Hence, both networks are operationally equivalent. \square

Since both semantics' are equivalent, we can choose the operational or denotational framework at our convenience. We usually derive the denotational semantics via operational computation paths, relying also on the fact that the semantics is schedule independent.

2.3 Compositionality

An important goal is to prove that the semantics of networks is conserved with respect to network composition as in Defs. 2.4 and 2.5. Knowing this, we can compose any two networks and be ensured that the resulting network carries out the intended computation. In other words, we need to prove Prop. 2.10 below – notice that a similar result exists for patterns [2], though the operations are of course defined differently. However, in order to do this we require a proper definition for composing mathematical objects of the form of Def. 2.8. The reasonable way to this is to gather types, inputs and outputs, eliminating those that are fed from one network into the other in case of sequential composition, exactly like we did for agents in Def. 2.2. Next, we need to combine quantum operations by tensoring with the identity map were necessary. As we shall see below, in this way we already recover most of Defs. 2.4 and 2.5. In fact, we only need to be careful in checking whether the quantum operations combine in the correct way, but, for this we can rely on the above-mentioned Prop. 2.10.

Proposition 2.10 *The semantics of networks is compositional, i.e.*

$$\llbracket \mathcal{N}_2 \cdot \mathcal{N}_1 \rrbracket = \llbracket \mathcal{N}_2 \rrbracket \cdot \llbracket \mathcal{N}_1 \rrbracket \quad (20)$$

$$\llbracket \mathcal{N}_1 \otimes \mathcal{N}_2 \rrbracket = \llbracket \mathcal{N}_2 \rrbracket \otimes \llbracket \mathcal{N}_1 \rrbracket \quad (21)$$

Proof. Suppose we have two networks $\mathcal{N}_1 = |_i \mathbf{A}_i(\mathbf{i}_{1,i}, \mathbf{o}_{1,i}) : Q_{1,i} \cdot \mathcal{E}_{1,i} \parallel \sigma_1$ and $\mathcal{N}_2 = |_i \mathbf{A}_i(\mathbf{i}_{2,i}, \mathbf{o}_{2,i}) : Q_{2,i} \cdot \mathcal{E}_{2,i} \parallel \sigma_2$ with semantics given by

$$\begin{aligned} \llbracket \mathcal{N}_1 \rrbracket &: \uplus_i Q_{1,i} \rightarrow \uplus_i Q'_{1,i} \cdot \mathbf{i}_1 \rightarrow \{(\mathbf{o}_1, \mathcal{L}_1^{\mathbf{o}_s})\} \\ \llbracket \mathcal{N}_2 \rrbracket &: \uplus_i Q_{2,i} \rightarrow \uplus_i Q'_{2,i} \cdot \mathbf{i}_2 \rightarrow \{(\mathbf{o}_2, \mathcal{L}_2^{\mathbf{o}_s})\}, \end{aligned} \quad (22)$$

with

$$\begin{aligned} \mathcal{L}_1 &: \mathcal{D}(\mathcal{H}_{I_1}) \rightarrow \mathcal{D}(\mathcal{H}_{O_1}) \otimes_i \rho_{1,i} \rightarrow \sum_j L_{1,j}(\sigma_1 \otimes_i \rho_{1,i}) L_{1,j}^\dagger \\ \mathcal{L}_2 &: \mathcal{D}(\mathcal{H}_{I_2}) \rightarrow \mathcal{D}(\mathcal{H}_{O_2}) \otimes_i \rho_{2,i} \rightarrow \sum_j L_{2,j}(\sigma_2 \otimes_i \rho_{2,i}) L_{2,j}^\dagger. \end{aligned} \quad (23)$$

We then find that

$$\llbracket \mathcal{N}_2 \rrbracket \cdot \llbracket \mathcal{N}_1 \rrbracket : \uplus_i Q_i \rightarrow \uplus_i Q'_i \cdot \mathbf{i} \rightarrow \{(\mathbf{o}, \mathcal{L}_2^{\mathbf{o}_2, s} \cdot (\mathcal{L}_1^{\mathbf{o}_1, s} \otimes \mathcal{I})), \forall \mathbf{o}_{1,s}, \mathbf{o}_{2,s}\} \quad (24)$$

where type, classical input and output are found by pointwise application of the rules in Eq.(1), and \mathcal{I} is the identity operation on $I_2 \setminus O_2$. Quantum operations in the above map states of the form $\otimes_i \rho_{1,i} \otimes_i \rho_{2,i}$ in $I = I_1 \uplus (I_2 \setminus O_2)$ to states in $O = O_2$, after tensoring them with $\sigma_1 \otimes \sigma_2$; On the other hand, the semantics of $\mathcal{N}_2 \circ \mathcal{N}_1$ is given by

$$\llbracket \mathcal{N}_2 \circ \mathcal{N}_1 \rrbracket : \uplus_i Q_i \rightarrow \uplus_i Q'_i \cdot \mathbf{i} \rightarrow \{(\mathbf{o}, [\mathcal{L}_2 \cdot (\mathcal{L}_1 \otimes \mathcal{I})]^{\mathbf{o}_s}), \forall \mathbf{o}_s\}, \quad (25)$$

where $\mathbf{o}_s = \mathbf{o}_{1,s} \uplus \mathbf{o}_{2,s}$, and the quantum operations operate on the same I and O as above. So we only need to check that first composing the quantum operations and then restricting them is the same as first restricting and then composing. However, this follows from the analogous result for ordinary patterns [2].

For parallel composition, consider $\mathcal{N}_2 = |_{i=1}^n \mathbf{B}_i(\mathbf{i}_{2,i}, \mathbf{o}_{2,i}) : Q_{2,i} \cdot \mathcal{E}_{2,i} \parallel \sigma_2$ instead, with semantics as above. We then find that

$$\llbracket \mathcal{N}_1 \rrbracket \otimes \llbracket \mathcal{N}_2 \rrbracket : \uplus_i Q_i \rightarrow \uplus_i Q'_i \cdot \mathbf{i} \rightarrow \{(\mathbf{o}, \mathcal{L}_1^{\mathbf{o}_1, s} \otimes \mathcal{L}_2^{\mathbf{o}_2, s}), \forall \mathbf{o}_{1,s}, \mathbf{o}_{2,s}\}, \quad (26)$$

where types, inputs and outputs are found by taking the (disjoint) union of those of the composing networks. Again, via the analogous result for ordinary patterns [2], we find that $(\mathcal{L}_1 \otimes \mathcal{L}_2)^{\mathbf{o}_s} = \mathcal{L}_1^{\mathbf{o}_1, s} \otimes \mathcal{L}_2^{\mathbf{o}_2, s}$, and therefore it follows that the above expression equals $\llbracket \mathcal{N}_1 \otimes \mathcal{N}_2 \rrbracket$. \square

2.4 Entanglement contexts

In the previous section, the networks \mathcal{N}_1 and \mathcal{N}_2 individually operate on tensor product states of the form $\otimes_i \rho_i$, in accordance with the fact that these are local inputs provided by each of the agents, as is made explicit in Def. 2.7. While this is sensible when considering a network operating in isolation – in which case entangled input states are specified as preparations – it is less

so when a network is only one factor in a complex compositional structure. Indeed, already for sequential composition the input state is in general no longer disentangled over agents, since it is fed in partly as output of a previous network computation, and there is no guarantee whatsoever that this output is a product state. Another subtle difference lies in the fact that input spaces of individual agents are combined when composing networks. Concretely, while each agent supplies local inputs ρ_1 when \mathcal{N}_1 is run separately, and ρ_2 when \mathcal{N}_2 is run separately, inputs to the composed network $\mathcal{N}_2 \circ \mathcal{N}_1$ are generally *not* of the form $\rho_1 \otimes \rho_2$. Yet another situation is that where agents have mixed state inputs because their inputs are entangled into a state on a larger system than that which the network operates on – in fact one view is that this is the *only* way in which mixed states arise. These are of course all typical manifestations of entangled states, but we want to stress the different situations in the context of distributed networks in which these arise. The point is, is entanglement preserved when applying operations to only part of the entangled state? We show below that it does, by proving the a quantum operation applied to the A -system of a state ρ existing on system AC does not touch the part of this state in C . Each of the aforementioned situations can be cast in this form, since A may be a system of several agents entangled with a group of other agents described by system C , as well as an input system of one agent only entangled with another input system C . The quantum operation itself is defined to map states on A to states on B , because we want to consider situations where a network's output spaces is different than its input space.

Proposition 2.11 *Suppose \mathcal{L} is a quantum operation on system A to system B such that*

$$\mathcal{L} : \mathcal{D}(\mathcal{H}_A) \longrightarrow \mathcal{D}(\mathcal{H}_B) : \rho_A \longrightarrow \rho_B = \sum_k L_k \rho_A L_k^\dagger. \quad (27)$$

Then for all quantum states ρ_{AC} living on a system AC , applying \mathcal{L} to one half of ρ_{AC} results in

$$\rho_{BC} = \sum_k (L_k \otimes I_C) \rho_{AC} (L_k^\dagger \otimes I_C) \quad (28)$$

Proof. Note first that any complex matrix can be written as a linear combination of Hermitian matrices, which in turn can be written as a sum of density matrices. Then by the spectral decomposition and linearity, it follows that for all complex matrices Z we have $\mathcal{L}(Z) = \sum_k L_k Z L_k^\dagger$. Writing $\rho_{AC} = \sum_{ijkl} \alpha_{ijkl} |i\rangle_A |j\rangle_C \langle k|_A \langle l|_C$, we find that

$$\begin{aligned} (\mathcal{N} \otimes I_C) \rho_{AC} &= \sum_{ijkl} \alpha_{ijkl} (\mathcal{N} \otimes I_C) (|i\rangle \langle k|^A \otimes |j\rangle \langle l|^C) \\ &= \sum_{ijkl} \alpha_{ijkl} \left(\sum_k (L_k |i\rangle \langle k|^A L_k^\dagger) \otimes |j\rangle \langle l|^C \right) \\ &= \sum_k (L_k \otimes I_C) \rho_{AC} (L_k^\dagger \otimes I_C), \end{aligned} \quad (29)$$

which proves the theorem. \square

The proof, while easy, is not trivial, and has several important consequences. First of all, it shows that our statement in Sec. 2.2 on the fact that networks are schedule-independent is true. Indeed, agents transform parts of a shared entangled state via local operations, and the above results shows that any order results in the same multilocal quantum operation. Next, while we have proved compositionality in the previous section considering only product state, we are now ensured that it holds also for arbitrary input states. We rely on Prop. 2.11 also in the next section, when we discuss sequencing several teleportation networks to transfer an entangled state from one agent to another.

3 Teleportation vs. quantum channels

In this section we prove that the teleportation protocol, described within our framework of networks of agents, is bisimilar to a direct quantum communication of the qubit to be teleported. We first give the network specification for direct quantum communication. Next, we give the network for the teleportation protocol, re-deriving the correctness of the protocol by developing its semantics in the local view elaborated in Sec. 2, which gives us a handle on qubit locations. A second goal of this section is to prove that it is bisimilar to a direct quantum channel. That is, by comparing the semantics of a direct quantum channel and that of teleportation – evolved in the correctness proof – we conclude that they do indeed define identical PTS's. By Prop. 2.11, this holds in arbitrary entanglement contexts. This is a nontrivial result because these agents may be entangled with the qubit to be teleported, and it is not a priori clear whether this entanglement is preserved throughout TP. This result has several consequences that we mention below. While the correctness of teleportation has been proved within other formal frameworks for distributed systems before [6,1], the bisimilarity approach, and specifically, taking into account arbitrary contexts, is new.

Consider a network of two agents named **A** and **B**. A direct quantum communication of a qubit is implemented simply by the network

$$\mathcal{N} = \mathbf{A} : \{1\}.(\mathbf{qc}!1) \mid \mathbf{B}.(\mathbf{qc}?1) \parallel \mathbf{0}, \quad (30)$$

where $\mathbf{0}$ is the null state. Note that **B** has an empty sort. The small-step semantics, given input $|\psi\rangle$, is derived in one step, and leads to the operational semantics

$$\llbracket \mathcal{N} \rrbracket : (\{1\}, -) \rightarrow (-, \{1\}).|\psi\rangle \Longrightarrow |\psi\rangle, \quad (31)$$

which immediately is the denotational semantics as well.

Consider the following network definition, which, as we derive explicitly below, implements the teleportation protocol.

$$\begin{aligned} \mathbf{a} &= \mathbf{A} : \{1, 2\}.[(\mathbf{c}!s_2s_1).M_{12}^{0,0}] \\ \mathbf{b} &= \mathbf{B} : \{3\}.[X_3^{x_2} Z_3^{x_1}.(\mathbf{c}?x_2x_1)] \\ \mathcal{N}_{TP} &= \mathbf{a} \mid \mathbf{b} \parallel E_{23} \end{aligned} \quad (32)$$

To derive the semantics of the above network, note that for the first step there is only one possibility, namely that agent **A** executes the local Bell measurement $M_{12}^{0,0}$. The latter requires a local quantum input from **A**, namely the qubit $|\psi\rangle$ that needs to be teleported. This is clearly the case since the pattern applies to qubit 1, which is not part of the system state E_{23} . So by rule (9) we need to apply the pattern to the first two qubits of $|\psi\rangle \otimes E_{23}$. Using first rule (9) and writing $\Gamma_{\mathbf{a}} = \emptyset[s_2s_1 \mapsto j_2j_1]$, we derive

$$\frac{|\psi\rangle_1 \otimes E_{23}, M_{12}^{0,0} \longrightarrow_{1/4} X^{j_2} Z^{j_1} |\psi\rangle_3, \Gamma_{\mathbf{a}}}{E_{23}, \emptyset, \mathbf{a} \Longrightarrow_{1/4} X^{j_2} Z^{j_1} |\psi\rangle_3, \Gamma_{\mathbf{a}}, \mathbf{A}.(\mathbf{c}!s_2s_1)}, \quad (33)$$

where for each of the values of j_2j_1 the transition occurs with the same probability of $1/4$. This reduction fires within the context of (12), which we do not write out explicitly. The next step is a classical rendez-vous between both agents (i.e. Alice calls Bob), as per rule (10). Defining $\Gamma_{\mathbf{b}} = \emptyset[x_2x_1 \mapsto j_2j_1]$, we get

$$\begin{aligned} X^{j_2} Z^{j_1} |\psi\rangle_3 \vdash (\Gamma_{\mathbf{a}}, \mathbf{A}.(\mathbf{c}!s_2s_1) \mid \emptyset, \mathbf{B} : \{3\}.[X_3^{x_2} Z_3^{x_1}.(\mathbf{c}?x_2x_1)]) \\ \Longrightarrow \Gamma_{\mathbf{a}}, \mathbf{A} \mid \Gamma_{\mathbf{b}}, \mathbf{B} : \{3\}.X_3^{x_2} Z_3^{x_1}. \end{aligned} \quad (34)$$

The last step of the computation is the execution of a local pattern by agent **B**, as follows.

$$X^{j_2} Z^{j_1} |\psi\rangle_3, \Gamma_{\mathbf{a}}, \mathbf{A} \mid \Gamma_{\mathbf{b}}, \mathbf{B} : \{3\}.X_3^{x_2} Z_3^{x_1} \Longrightarrow |\psi\rangle_3, \Gamma_{\mathbf{a}}, \mathbf{A} \mid \Gamma_{\mathbf{b}}, \mathbf{B} : \{3\} \quad (35)$$

The only probabilistic transition is the first one, due to the Bell measurement. However, we see that the four branches lead to identical final system state and agents specifications. Furthermore, since there is no classical input or output, we can trace out the different local states. Thus, adding the probabilities as specified in Def. 2.7, we find that for any input $|\psi\rangle$, we have

$$\llbracket \mathcal{N} \rrbracket_{TP} : (\{1, 2\}, \{3\}) \rightarrow (-, \{3\}).|\psi\rangle \Longrightarrow |\psi\rangle \quad (36)$$

In other words, we find that $\llbracket \mathcal{N}_{TP} \rrbracket \equiv \llbracket \mathcal{N} \rrbracket$.

Note that by linearity, the above derivation also works for mixed states.

We have shown that a direct quantum channel is operationally equivalent to the teleportation protocol. Suppose however, that two agents wish to exchange a qubit whilst they are contained in a larger network of agents. Can we say anything about the equivalence of both procedures in this context? By Prop. 2.11, we know that the entanglement with the larger system is conserved. Indeed, since teleportation just implements an identity channel, applying teleportation to one half of the mixed state ρ_{AC} results in the state ρ_{BC} . Furthermore, suppose agent **A** wants to send an n -qubit entangled state to **B**. Then by the same result, **A** can just apply the teleportation protocol n times, and, since entanglement is conserved, the state is transferred unchanged. It remains to be seen whether the conservation of correlations

between agents can also be employed in higher-level applications. Considering the fact that shared entanglement provides much of the extra power in distributed quantum systems, this behavior seems a promising primitive.

4 Conclusion

In this paper, we develop a formal model for distributed measurement-based quantum computations. We adopt an agent-based view, such that computations are described locally where possible. Because the network quantum state is in general entangled, we need to model it as a global structure, reminiscent of global memory in classical agent systems. Local quantum computations are described as measurement patterns. Since measurement-based quantum computation is inherently distributed, this allows us to extend naturally several concepts of the measurement calculus (MC) [2], a formal model for such computations. Just as in MC, we aim at defining an *assembly language*, i.e. we assume that computations are well-defined and do not concern ourselves with verification techniques. The operational semantics for systems of agents is given by a probabilistic transition system, and we define operational equivalence in a way that it corresponds to the notion of *bisimilarity*. The denotational semantics is given by a set of quantum operations, together with type information which determines the localization of resources. Both forms of semantics are proved to be equivalent, and we define a notion of network composition such that the semantics is preserved with respect to this operation. Moreover, we show that within the larger entanglement contexts, the semantics is also conserved. With this in place, we prove that teleportation is bisimilar to a direct quantum channel, and, by the above-mentioned result, this also holds within the context of a larger network. Though the proof is quite simple, it is important within the context of agent systems. Indeed, the possibility of inheriting agent correlations via teleportation means that, for example, collaborating agents are not cut off from each other when part of the shared data is transferred. Rather, the correlations are preserved in an oblivious manner. That this is the case if qubits are transported physically is clear, but that this remains so even if a general protocol is employed is maybe more surprising. It remains to be investigated how this preservation of the entanglement context can be exploited in more general situations. Also, there are other interesting situations to be investigated, such as the so-called *channel inequalities*, and more elaborate communication protocols. This is the subject of current investigations.

References

- [1] Abramsky, S. and B. Coecke, *A categorical semantics of quantum protocols*, Technical Report PRG-RR-04-02, Oxford University Computing Laboratory (2004).

- [2] Danos, V., E. Kashefi and P. Panangaden, *The measurement calculus* (2004).
URL <http://arxiv.org/abs/quant-ph/0412135>.
- [3] Danos, V., E. Kashefi and P. Panangaden, *1-qubit versus 2-qubit measurement-based quantum computation* (2005).
- [4] D'Hondt, E., "Distributed measurement-based quantum computation," Ph.D. thesis, Vrije Universiteit Brussel (2005), (in preparation).
- [5] D'Hondt, E. and P. Panangaden, *Quantum weakest preconditions*, in: P. Selinger, editor, *Proceedings of the 2nd Workshop on Quantum Programming Languages (QPL04)*, Turku Centre for Computer Science (2004).
URL <http://arxiv.org/abs/quant-ph/0501157>.
- [6] Gay, S. J. and R. Nagarajan, *Communicating quantum processes*, in: P. Selinger, editor, *Proceedings of the 2nd Workshop on Quantum Programming Languages (QPL04)*, Turku Centre for Computer Science (2004).
URL <http://arxiv.org/abs/quant-ph/0409052>.
- [7] Lalire, M. and P. Jorrand, *A process algebraic approach to concurrent and distributed quantum computation: operational semantics*, in: P. Selinger, editor, *Proceedings of the 2nd Workshop on Quantum Programming Languages (QPL04)*, Turku Centre for Computer Science (2004).
URL <http://arxiv.org/abs/quant-ph/0407005>.
- [8] Leung, D. W., *Quantum computation by measurements*, Int. J. of Quant. Comp. **2** (2004), pp. 33–43.
URL <http://arxiv.org/abs/quant-ph/0310189>.
- [9] Nielsen, M. A., *Universal quantum computation using only projective measurement, quantum memory, and preparation of the 0 state*, Phys. Lett. A **2-3** (2003), pp. 96–100.
URL <http://arxiv.org/abs/quant-ph/0108020>.
- [10] Nielsen, M. A. and I. Chuang, "Quantum computation and quantum information," Cambridge university press, 2000.
- [11] Raussendorf, R., D. E. Browne and H. J. Briegel, *Measurement-based quantum computation on cluster states*, Phys. Rev. A **68** (2003), p. 022312.
URL <http://arxiv.org/abs/quant-ph/0301052>.
- [12] Walther, P., K. J. Resch, T. Rudolph, E. Schenck, H. Weinfurter, V. Vedral, M. Aspelmeyer and A. Zeilinger, *Experimental one-way quantum computing*, Nature **434** (2005), pp. 169–176.
URL <http://arxiv.org/abs/quant-ph/0503126>.