International Journal of Theoretical Physics, Vol. 42, No. 9, September 2003 (© 2003)

Discrete Quantum Causal Dynamics

Richard F. Blute,[1,3] Ivan T. Ivanov,[1] and Prakash Panangaden[2]

Received July 16, 2002

We give a mathematical framework to describe the evolution of open quantum systems subject to finitely many interactions with classical apparatuses and with each other. The systems in question may be composed of distinct, spatially separated subsystems which evolve independently, but may also interact. This evolution, driven both by unitary operators and measurements, is coded in a mathematical structure in such a way that the crucial properties of causality, covariance, and entanglement are faithfully represented. The key to this scheme is the use of a special family of spacelike slices—we call them *locative*—that are not so large as to result in acausal influences but large enough to capture nonlocal correlations.

KEY WORDS: discrete quantum systems; causality; entanglement.

1. INTRODUCTION

We propose a uniform scheme for describing a quantum system interacting with a network of classical objects. The system in question may be composed of distinct spatially separated subsystems that evolve independently, but may also interact with each other at various points as well as with the classical objects. When analyzing physical laboratory experiments on quantum systems, we frequently abstract away from the concrete experimental setup and from the particular details of the machinery involved. What we usually keep is the description of the quantum system—and its spatially separated subsystems—in terms of wave functions or density matrices and unitary operators as well as the changes of the quantum system induced by the interactions with classical devices. Crucial properties of the evolution such as the causal ordering, covariance of the description for different observers, and quantum entanglement between distinct subsystems should be completely reflected in any such description.

[1] Department of Mathematics and Statistics, University of Ottawa, Ottawa, Ontario, Canada.
[2] School of Computer Science, McGill University, Montreal, Quebec, Canada.
[3] To whom correspondence should be addressed at Department of Mathematics and Statistics, University of Ottawa, 585 King Edward Avenue, Ottawa, Ontario, Canada, K1n 6N5; e-mail: rblute@mathstat.uottawa.ca.

0020-7748/03/0900-2025/0 © 2003 Plenum Publishing Corporation

The basis of our representation is the graph of events and causal links between them. An event could be one of the following: a unitary evolution of some subsystem, an interaction of a subsystem with a classical device (a measurement), or perhaps just the coming together or splitting apart of several spatially separated subsystems. Events will be depicted as vertices of a directed graph. The edges of the graph will represent propagation of quantum systems between two different events. The vertices of the graph are then naturally labelled with operators representing the corresponding interactions.

Of course, the processes of unitary evolution and measurement take a certain amount of time; but we are only interested in the causal relations between such events and this allows us to consider them as point-like vertices on the graph. Thus we are thinking of the duration between events as being longer than the duration of any event so that no causal information is lost when we represent interactions as pointlike events.

The structure described thus far reflects the kinematic properties of the quantum system. To describe the dynamics we need a composition of the operators assigned to the vertices of the graph. This composition is described in terms of the combinatorics of the underlying graph. Causal relations are made explicit and we prove that no influences breaking causality arise in our scheme. The possible entanglement between spatially separated subsystems—represented by distinct edges of the graph—is also accounted for. Thus, our framework allows one to represent locality of interaction—i.e. causal influences do not propagate outside the causal "cone"—while allowing the expression of nonlocal correlations which occur when one has quantum entanglement. The tension between causal evolution and quantum entanglement is resolved.

1.1. Relation to Other Work

Next we outline the relations of our proposal to some recent approaches to quantum mechanics and quantum gravity.

1.1.1. Decoherent Histories

The goal of quantum mechanics is to determine the probability of an event or a sequence of events, thus one might hope to assign probabilities to the histories of the quantum system. In order for the probabilities to be additive in the usual sense, the histories have to be mutually noninterfering. Sets of histories obeying this condition are selected with the use of a special bilinear form on histories—the decoherence functional (Gell-Mann and Hartle, 1993; Griffiths, 1996; Omnés, 1994).

A particular history is mathematically represented as a linearly ordered sequence of projection operators in the Hilbert space of the quantum mechanical

system. But the linear causal ordering of the events in a history is too restrictive in many experimental situations, in particular when analyzing spatially separated entangled quantum systems. This issue is even more pressing for quantum cosmology considerations. An application of the histories approach to quantum field theory on a curved space-time (Blencowe, 1991) must assume the existence of a globally hyperbolic manifold, and thus via the associated foliation, a linear ordering of the histories of the quantum field.

Our proposal for describing the evolution of an open quantum system can be considered as describing a single history in a set of histories. The important point is that events are no longer linearly ordered by temporal order but, rather, partially ordered with respect to the causal order. This allows one to capture the notion of causal evolution in a manifestly covariant fashion. The decoherence condition for histories has an immediate generalization for histories described by more general graphs as proposed here.

1.1.2. Causal Sets

Causal sets form the basis of an approach to quantum gravity mainly advocated by R. Sorkin and collaborators (Bombelli *et al.*, 1987; Sorkin, 1991), where the basic idea is to take the notion of causality as the primitive. In classical relativity, the structure of the space-time manifold together with a metric of Lorentzian signature determines the causality relation. An important observation is that the causal structure is conformally invariant, i.e. determined by only the conformal equivalence class of the metric and hence more primitive than the metric. Various proposals for quantum gravity—for example, the twistor program (Penrose and MacCallum, 1972)—have taken as their point of departure the idea that the causal structure is more fundamental than the metric structure.

In the causal sets approach, one takes the point of view that, at the smallest length scales, space-time is inherently discrete and that the causal structure, the "light cones," are fundamental. This leads naturally to the idea of a partially ordered set (poset for short) where the elements are events and two events are related by causality. The main interest is in approximating continuous space-times with such structures and defining processes that would generate these structures, with a view to an eventual theory of quantum gravity. Though the aims are rather different the issues connected with causality are closely related.

Causal sets are further motivated by the idea that a discrete structure would avoid the singularities that plague physics (both classical and quantum). The assumption that space-time should be a continuous manifold is one of the ingredients that leads to the problematic singularities of quantum field theory and general relativity. In the causal sets approach, space-time is a discrete structure, thus possibly avoiding these singularities, the idea being that at the Planck scale, continuous geometry gives way to discrete geometry.

One way to think of this is that one approximates a manifold as one "sprinkles" more and more points into the causal set in a uniform fashion. Applications and extensions of these ideas can be found in papers such as (Markopoulou, 2000; Markopoulou and Smolin, 1997; Raptis, 2000), although this list is by no means exhaustive. For us, a finite causal set is the kinematic framework on which we describe evolution and information flow.

1.1.3. Quantum Causal Histories

The notion of *quantum causal history* was introduced by Markopoulou in (Markopoulou, 2000). One begins with a poset (causal set) and assigns Hilbert spaces to the vertices and evolution operators to sets of edges. However, within this framework, one is quickly led to violations of causality—as the author herself notes—essentially because the slices used are "too global." She mentions the possibility of working with a dual view. In fact, in our work, we take such a dualized view as our starting point. In other words we assign operators representing evolution or measurement to vertices and Hilbert spaces to the edges. However, if we only work locally we get a causal theory but lose the possibility of capturing nonlocal correlations.

1.2. The Importance of Graphs

A graph can be seen as a generalization of a poset in the following sense. A poset merely records that an element x precedes y. In a poset when one writes $x \leq y$ then, depending on the context, one is stating something like x causally precedes y, x implies y, or any of several other possibilities. On the other hand, a graph keeps track of the different ways in which x might precede y.

In this work, we are particularly interested in modelling the idea that information can flow from one event to another in a number of different ways, *along different paths or channels*. We would like to keep track of all these various independent paths. The structure of a poset is inadequate for achieving this, as we would like to say that x causally precedes y in several different ways. This naturally suggests that we pass from posets to more general graphs.

Many recent experiments feature spatially distributed quantum systems. When entangled quantum subsystems come back together in the same space-time region, the description of the resulting system is causally influenced by all events in the paths of the subsystems. In particular, a past event could influence the future events in several distinct ways through different paths. Our scheme is well adapted for analyzing experiments featuring spatially separated quantum entangled entities and could be used in the field of quantum information processing to analyze information flow situations.

1.3. Contents of the Present Paper

Section 2 presents the basic ideas of our scheme via an example. Section 3 discusses the basic physical ideas involved. In the first subsection we review the notions of measurements and interventions. In the next subsection we give the dynamical prescription in a special case and in the final subsections we give the general prescription and prove covariance.

2. CAUSAL INFORMATION FLOW VIA EXAMPLES

Consider a quantum system evolving in space-time while being subjected to interactions with classical observers at a number of points. The causal and spatiotemporal relations in the system will be represented by a directed acyclic graph (hereafter called a *dag*). The vertices of the graph—which will be drawn as boxes—represent the events in the evolution of the system. An event could be a measurement by a classical observer, a local unitary evolution, or just a splitting of a subsystem into several spatially separated subsystems, which however could still share an entangled common state. The propagation of the different subsystems will be indicated by the edges of the graph.

There are a number of causal relations between edges and vertices. A vertex v_1 is said to *immediately precede* v_2 if there is a (directed) edge from v_1 to v_2. We write $v_1 \preceq v_2$ for the reflexive transitive closure of immediate precedence; thus $v \preceq v$ always holds and $v_1 \preceq v_2$ means that there is a *directed path* from v_1 to v_2 (possibly of length zero). When $v_1 \preceq v_2$ we sometimes say v_1 is "to the past of" v_2 and dually "v_2 is to the future of v_1." When we draw a poset we typically leave out the self-loops and only draw the minimal number of edges needed to infer all the others; the so-called "Hasse diagram" of the poset. We note that our graphs will have initial and final "half-edges," i.e. edges with only one endpoint. Physically we have some quantum states incoming (or "prepared") followed by some interactions and some outgoing state.

The relation between vertices induces a causal relation between edges. We say that an edge e_1 is to the past of another edge e_2 if the terminal vertex of e_1, say v_1 and the initial vertex of e_2, say v_2, satisfy $v_1 \preceq v_2$. Note that we could have $v_1 = v_2$. An initial edge is not to the future of any edge, nor is a final edge to the past of any other edge. If two edges are not causally related, we say that they are "spacelike separated" or acausal. Note that two spacelike separated edges could share a common terminal vertex or a common initial vertex (but since we have a graph, not both). A *space-like slice* is defined as a set of pairwise acausal edges. Henceforth, whenever we say "slice" we will always mean "spacelike slice." Note that the initial (or final) edges form a spacelike slice. We call this the *initial (final) slice*.

For example, for the graph of Fig. 1, the set of edges $\{e_c, e_d, e_e\}$ form a spacelike slice. Another example is the set $\{e_f, e_d, e_e\}$. The edges e_a and e_b form

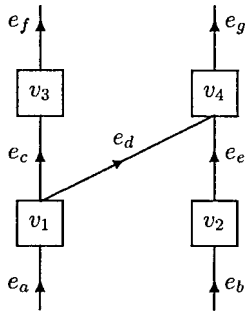

Fig. 1.

the initial slice. The edges e_a, e_b, e_f, and e_g are half-edges, with e_a and e_b initial, and e_f and e_g final.

Associated with any edge e_i is an observer who has access to a subsystem of the complete quantum system. Thus the edges represent local information. Each edge e_i is assigned a density matrix ρi in a Hilbert space $\mathcal{H}_i$.[4] The density matrix ρi describes the knowledge about the quantum system available to the local observer at the edge e_i. More generally, density matrices will be associated to spacelike slices. For a spacelike slice consisting of edges $\{e_{i_1}, \ldots e_{i_p}\}$, the assigned density matrix will be denoted $\rho_{i_1}, \ldots i_p$. This density matrix describes the subsystem of the whole quantum system for that spacelike slice. Every spacelike slice has also a Hilbert space which is the tensor product of the Hilbert spaces of the edges forming the slice. However, the density matrix associated with the slice is not in general a tensor product of the density matrices on the edges. If it were, we could not capture nonlocal quantum correlations.

The graph of Fig. 1, represents a quantum system Q which starts evolving from a state in which Q consists of two spatially separated subsystems Q_a and Q_b described by density matrices ρ_a and ρ_b, respectively, in Hilbert spaces $\mathcal{H}_a$ and $\mathcal{H}_b$. The initial edges e_a and e_b form the initial slice in this simple system. We will follow the convention that if the initial slice consists of several edges, the initial state of the whole system is a tensor product state, i.e. the subsystems are not entangled. For the above example, $\psi_{\text{init}} = \psi_a \otimes \psi_b$ and $\rho_a = |\psi_a\rangle\langle\psi_a|$ and $\rho_b = |\psi_b\rangle\langle\psi_b|$. Entangled subsystems on distinct edges will always have at least one event in the common past. Thus we always explicitly represent the interaction which caused the entanglement.

Each vertex v_i of the graph is labelled with an operator T_i which describes the process taking place at the corresponding event. The operator T_i at a given event v_i takes density matrices on the tensor product of Hilbert spaces living on the incoming edges at v_i to density matrices on the tensor product Hilbert space of

[4] Throughout the paper, we assume that the graph and the dimensions of all Hilbert spaces are finite.

outgoing edges. The process at a vertex could be an *intervention*[5] corresponding to a positive operator-valued measure (POVM) (Nielsen and Chuang, 2000; Peres, 1995) or a unitary transformation. Or instead of an external or unitary action there could be several quantum subsystems that come together and then split apart, possibly in a different way. We will consider this last case as a particular instance of a unitary evolution with identity evolution operator. As a simple example, in the case of an event corresponding to unitary evolution by a unitary operator U, we have the usual expression:

$$\rho^{\text{in}} \mapsto \rho^{\text{out}} = U\rho^{\text{in}}U^{\dagger} \tag{1}$$

The general expression for an operator associated to an event will be discussed fully in the next section, see Eq. (3).

Here we will discuss some of the conditions such a dynamical scheme has to satisfy in order to reflect causality and other physical properties of the quantum system. Causality is the condition that the density matrix on a given edge should not depend on the actions performed at vertices which are acausal to this edge or are in its future. For example, referring back to Fig. 1, we would like any quantum evolution rule to say that the density matrix at e_g is unaffected by the intervention at v_3 or the density matrix at e_f is unaffected by the intervention at v_2. A general unitary evolution between the states of two spacelike slices is easily shown to violate this condition. Therefore we need to incorporate some sort of locality condition into the evolution scheme.

It is not hard to formulate such an evolution scheme. For example, one could work with the dual picture and have evolution occur along edges with density matrices at the vertices. It is not hard to formulate rules which would enforce causality properly in such a framework. Unfortunately this rules out quantum correlations across spatially separated subsystems. Thus, the evolution scheme cannot be too local because entangled subsystems of the quantum system could separate and later come together at a vertex.

Consider the system shown in Fig. 2. The quantum system represented in this graph is as follows. The system is prepared in a state ψ_a as indicated by the density matrix $\rho_a = |\psi_a\rangle\langle\psi_a|$ on the incoming edge. At the vertex v_1 the system splits into two spatially separated subsystems on the edges e_b and e_c which, in general, are still described by a global entangled state. The local transformations T_2 and T_3 will, in general, preserve the entanglement and the global state will be still entangled on the spacelike slice $\{e_d, e_e\}$. The two subsystems come together at the vertex v_4. The two local density matrices ρ_d and ρ_e are not sufficient to reconstruct the entangled state of the system described by ρ_f. The off-diagonal terms of ρ_f are not reflected in the local density matrices, ρ_d and ρ_e. We need to include information

[5] Interventions are generalized measurements where a quantum subsystem could be discarded (Peres, 2000a). This will be discussed more fully below.

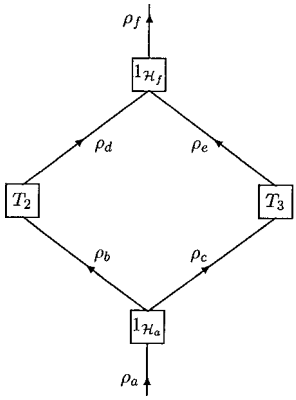

Fig. 2.

about the history of the state on the spacelike slice $\{e_d, e_e\}$ in order to reconstruct the global state. One possibility is to work with global spacelike slices, and show that the scheme is generally covariant in the sense of being slice-independent. In our approach, certain preferred (not necessarily global) spacelike slices account for all entanglement.

The rules for constructing and labeling the graphs given so far reflect the kinematics of the quantum system. Specifying the dynamics amounts to a prescription for how to obtain the density matrices on every edge from the density matrix on the initial slice and the operators at the vertices of the graph. This prescription will be given below in the next section.

3. DYNAMICS ON GRAPHS

3.1. Measurements and Interventions

For standard material on density matrices, positive operator-valued measures (POVMs), and completely positive operators, we refer to (Davies, 1976; Nielsen and Chuang, 2000; Preskill, . . .).

The *measurement* of a property of a quantum system involves interaction with a classical apparatus. When a classical apparatus measures an observable of a quantum subsystem sitting inside a larger system the appropriate mathematical formalism for such generalized measurement is that of POVM. Let the possible outcomes of the measurement be labelled by the letter $\mu \in \{1 \ldots N\}$. The measurement is described by a unitary interaction between the apparatus and the quantum system. The classical apparatus has a preferred basis of states indexed by μ. After the measurement, the apparatus appears in one of these preferred states. Since we are only interested in describing our quantum subsystem, we trace out all the remaining degrees of freedom. Effectively to every outcome μ is

associated an operator F_μ. Consider the family of positive operators $E_\mu = F_\mu^\dagger F_\mu$. For a generalized measurement the family must satisfy the condition $\sum_\mu E_\mu = I$. The probability p_μ for obtaining a measurement result labelled by μ is then given by: $p_\mu = T_r(E_\mu \rho) = T_r(F_\mu \rho F_\mu^\dagger)$. The density matrix of Q_1 after the measurement with outcome μ is given by

$$\tilde{\rho}_\mu = \frac{1}{p_\mu} F_\mu \rho F_\mu^\dagger \tag{2}$$

where ρ is the density matrix before the measurement and p_μ is inserted to normalize the resulting density matrix to unit trace.

A more convenient way of using the density matrix formalism is as follows. We want to label the states resulting from measurements with the probabilities with which these states appear. These states with probabilities are described by unnormalized density matrices, i.e. positive, self-adjoint operators with trace less than or equal to 1. The above formula for $\tilde{\rho}_\mu$ can be written in terms of unnormalized density matrices as follows: $\rho'_\mu = F_\mu \rho F_\mu^\dagger$. The trace of this density matrix is precisely p_μ, the probability that the outcome labelled by μ was realized. This makes sense even if p_μ is zero. Henceforth, whenever we talk about density matrices we will mean these unnormalized density matrices, which describe states with a probability attached to them. If $\mathcal{H}$ is any Hilbert space, then the set of all such density matrices will be denoted DM($\mathcal{H}$).

Even more general measurement processes could be considered if the observer discards part of the quantum system during the process of measurement. The appropriate mathematical formalism for describing these generalized measurements is that of *intervention operators* (Peres, 2000a). In the process of measurement, the density matrix changes according to:

$$\rho'_\mu = \sum_m A_{\mu m} \rho A_{\mu m}^\dagger \tag{3}$$

The families of maps $A_{\mu m}$ now act in general from one Hilbert space to another, i.e. for fixed μ and m they correspond to rectangular matrices. The label μ again distinguishes the set of possible outcomes and the letter m labels the degrees of freedom discarded during this generalized measurement. Since the maps $A_{\mu m}$ come from measurements realized by unitary operator on some larger Hilbert space they again satisfy a completeness condition: $\sum_{\mu m} A_{\mu m}^\dagger A_{\mu m} = I$, where I is the identity operator in the appropriate Hilbert space. Notice that if the labels μ and m are absent in (3) the equation describes unitary evolution. Since the events we consider are generalized measurements or unitary evolutions, Eq. (3) is the appropriate mathematical representation of those processes in full generality. Equation (3) has precisely the form of a general completely positive linear map (Davies, 1976) and we will use here also the terminology *intervention operators* for operators on density matrices of this form.

3.2. The Dynamical Prescription

We are now ready to start discussing the dynamics of a quantum system represented by a dag G. Dynamics will be described by supposing that we are given a density matrix on the initial spacelike slice, and then giving a prescription for calculating the density matrices of future spacelike slices. In essence, we are propagating the initial data throughout the system.

To each vertex $i \in G$ will be assigned an intervention operator T_i, and to each edge e_j will be assigned a Hilbert space $\mathcal{H}_j$. We note that all incoming (or outgoing) edges of a given vertex are pairwise acausal and thus form a spacelike slice. Thus there will be a density matrix ρ_i^{in} associated to the slice of the incoming edges. Then one obtains the density matrix for the slice of the outgoing edges by:

$$\rho_i^{\text{in}} = T_i(\rho_i^{\text{out}}).$$

Notice that more generally, for two acausal vertices, the sets of incoming or outgoing edges are pairwise acausal. Thus, the associated intervention operators will act on different Hilbert spaces and hence commute.

We begin with an illustrative example. Consider the dag of Fig. 3. Given the state on the initial slice, the operators at the events propagate the state to the future. In the example of Fig. 3 we have: $\rho_c = T_1(\rho_a)$, $\rho_{fde} = T_2(\rho_b)$. However, the next intervention operator T_3 must act on the so far undefined density matrix ρ_{cd}. T_3 takes density matrices on $\mathcal{H}_c \otimes \mathcal{H}_d$ to those on $\mathcal{H}_g \otimes \mathcal{H}_d$. By extending T_3 with the appropriate identity operators, we can view it as a map from DM($\mathcal{H}_c \otimes \mathcal{H}_d \otimes \mathcal{H}_e \otimes \mathcal{H}_f$) to DM($\mathcal{H}_e \otimes \mathcal{H}_f \otimes \mathcal{H}_g \otimes \mathcal{H}_h$). Then we can define the density matrix on another spacelike slice, namely $\rho_{fghe} = T_3(\rho_c \otimes \rho_{fde})$. Similarly $\rho_{fdi} = T_4(\rho_{fde})$ and so on. Starting from density matrices on the initial edges and using the intervention operators associated with the vertices—extended with identities as needed—we obtain density matrices on specific spacelike slices.

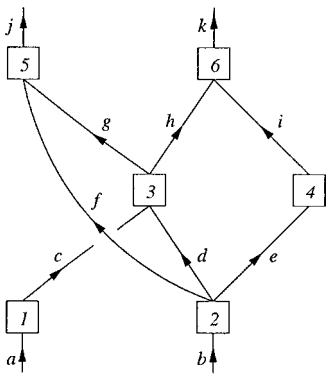

Fig. 3.

The above inductive process for propagating density matrices can be applied to any system described by a dag. However, the procedure only gives the density matrices for certain spacelike slices within the dag. For example, this procedure does not yet yield a matrix for the slice *de*. To calculate such density matrices, we would also have to make use of the trace operator. Before extending the procedure to such slices, we first consider those for which the above process is sufficient. We call these slices *locative*.

Definition 3.1. Let G be a dag, and L a slice of G. Consider the set of all vertices V which are to the past of some edge in L. Let I be the set of initial edges in the past of L. Consider all paths of maximal length beginning at an element of I and only going through vertices of V. Then L is *locative* if all such paths end with an edge in L.

In our example, the locative slices are the following:

$$a, b, ab, c, cb, def, adef, cdef, efgh,$$

$$adfi, cdfi, fghe, fghi, fgk, hej, hij, jk$$

while, for example, *de* is not locative. Note that the fact that maximal slices are always locative follows immediately from the definition of locative.

We now describe the general rule for calculating the density matrices on locative slices. Associated with each locative slice L is the set I of initial edges in the past of L. We choose a family of slices that begins with I and ends with L in the following way. Consider the set of vertices V between the edges in I and the edges in L. Because L is locative we know that propagating slices forwards through the vertices in V will reproduce L. Let $M \in V$ be such that the vertices in M are minimal in V with respect to causal ordering. We choose arbitrarily any vertex u in M, remove the incoming edges of u, and add the outgoing edges of u to the set I obtaining a new set of edges I_1. It is clear that I_1 is spacelike and locative. Proceeding inductively in this fashion we obtain a sequences of slices $I = I_0, I_1, I_2, \ldots I_n = L$, where, n is the cardinality of V. Of course, this family of slices is far from unique.

The dynamics is obtained as follows. Recall that the states on initial edges are assumed not to be entangled with each other so that one can obtain the density matrix on any set of initial edges, in particular I, by a tensor product. Let ρ_0 be the density matrix on I. We look at the vertex u that was used to go from I to I_1 and apply the intervention operator T assigned to this vertex—possibly augmented with identity operators as in the example above. Proceeding inductively along the family of slices, we obtain the density matrix ρ_n on L.

The important point now is that ρ_n does not depend on the choice of slicing used in going from I to L. This can be argued as follows. Suppose we have a locative slice S and two vertices u and v which are both causally minimal above

S and acausal with respect to each other. Then we have four slices to consider, S, S_u, S_v, and S_{uv} where by S_u we mean the slice obtained from S by removing the incoming edges of u and adding the outgoing edges of u to S and similarly for the others. It is clear, in this case, that the intervention operators assigned to u and to v commute and the density matrix computed on S_{uv} is independent of whether we evolved along the sequence $S \to S_u \to S_{uv}$ or $S \to S_v \to S_{uv}$. Now when we constructed our slices at each stage we had the choice between different minimal vertices to add to the slice. But such vertices are clearly pairwise acausal and hence, by the previous argument applied inductively, the evolution prescription is independent of all possible choices. We summarize this argument in the following proposition

Proposition 3.2. (*Covariance*)**.** *The density matrix on a locative slice is independent of the family of slices used to compute it.*

So far we have defined density matrices on locative slices only. To define density matrices on general spacelike slices we will need to consider partial tracing operations.

3.3. General Slices

Recall that when one has subsystems Q_1 and Q_2 of a quantum system Q, the Hilbert space for Q may be decomposed as $\mathcal{H}_1 \otimes \mathcal{H}_2$ where $\mathcal{H}_i$ represents Q_i. The density matrix for Q_1 is obtained by tracing over $\mathcal{H}_2$. To obtain the density matrix of a general spacelike slice S we proceed as follows. First note that there is a natural partial ordering[6] on slices (recall that we always mean "spacelike slices" whenever we say "slices") given as follows

$$S_1 \sqsubseteq S_2 \text{ if } \forall_x \in S_1 \exists_y \in S_2 \cdot x \preceq y \text{ and } \forall_y \in S_2 \exists_x \in S_1 \cdot \preceq y$$

where $x \preceq y$ means that x causally precedes y. The collection of slices forms a finite poset. We find the *minimal* (in the order $\sqsubseteq$) locative slice M that contains S and trace over the Hilbert spaces on edges in $M \setminus S$. Such a locative slice M always exists because maximal spacelike slices are always locative.

Lemma 3.3. *Given a slice S, there is a unique minimal slice which contains S.*

Proof: We claim that the locative slice obtained by evolving through those vertices to the causal past of S is the unique minimal locative slice. More precisely, let $V = \{v_1, \ldots, v_p\}$ be the set of vertices to the past of S. Let $I = \{e_1, \ldots, e_q\}$

[6] This ordering is well-known in computation theory where it is called the Egli–Milner order (Plotkin, 1976).

be the set of initial edges to the past of S. Constructing a sequence of slices by incrementally incorporating the vertices of V in a manner similar to what we did in the previous subsection, we get a locative slice M containing S.

First, note that M is locative by definition and also it contains S. Second, note that any slice containing S (in particular S itself) must be to the future of all these vertices. Thus M is indeed minimal since it lies to the future of the fewest vertices. □

Definition 3.4. We shall refer to M as the *least locative slice* of the edge S.

Starting with the density matrices on the edges of I and applying the operators associated with the vertices of V, we obtain the density matrix on the locative slice M. Suppose that S consists of edges $\{e_1, \ldots e_i\}$ Let M—the minimal locative slice through S—consist of edges $\{e_1, \ldots, e_i, e'_1, \ldots, e'_j\}$. The density matrix $\rho_{1,\ldots,i,1',\ldots,j'}$ on M is an element of the space $End(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_i \otimes \mathcal{H}'_1 \otimes \cdots \otimes \mathcal{H}'_j)$. Let $Tr^{1'\ldots j'}$ be the partial trace operation

$$End(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_i \otimes \mathcal{H}'_1 \otimes \cdots \otimes \mathcal{H}'_j) \rightarrow End(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}'_i)$$

which traces over the primed Hilbert spaces.

Definition 3.5. (Density matrix associated with a slice). The density matrix ρ at the slice S is defined to be:

$$\rho_S = Tr^{1'\ldots j'}(\rho_{1,\ldots,i,1',\ldots,j'}). \tag{4}$$

Of course if M and S coincide then no tracing is done.

Proposition 3.6. *The prescription for computing density matrices is causal in the sense that the only vertices to the past of S can affect the density matrix on S.*

Proof: This is immediate from the construction and uniqueness of the minimal locative slice. The explicit construction of the minimal locative slice shows that only intervention operators to the past of S can affect the density matrix on S. □

For ease of presentation the rest of the discussion will focus on the case where S is a single edge. In general, an edge e_i is contained in many locative slices and we could just as well have defined ρ_i by tracing over the complementary degrees of freedom in any of these locative slices. However we have to compute the density matrices differently in order to ensure causality. To clarify the discussion consider the quantum system represented by the graph on Fig. 4.

Let the initial ρ_a be the density matrix of a maximally entangled state of two spin 1/2 subsystems: $\rho_a = |\psi_a\rangle\langle\psi_a|$, where $\psi_a = 1/\sqrt{2}(\psi_1^{\uparrow} \otimes \psi_2^{\uparrow} + \psi_1^{\downarrow} \otimes \psi_2^{\downarrow})$. At the first vertex the two subsystems separate with no classical intervention.

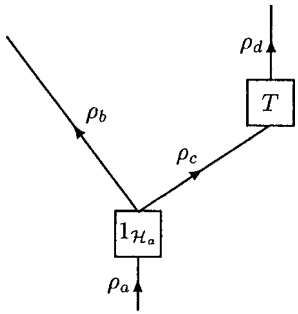

Fig. 4.

Therefore $\rho_{bc} = \rho_a$. The slice $\{e_b, e_c\}$ is the least locative slice for the edge e_b and we can compute the density matrix associated to this edge: $\rho_b = Tr^c \rho_{bc} = 1/2(|\psi_1^\uparrow\rangle\langle\psi_1^\uparrow| + |\psi_1^\downarrow\rangle\langle\psi_1^\downarrow|)$. Next, let the intervention at the second vertex be a measurement on the corresponding subsystem with the result that the spin was found to be in the state $\psi_2^\uparrow$. The intervention operator is the projection operator on this state of the second subsystem: $T(\rho) = 2P_2^\uparrow \rho P_2^\uparrow$. We obtain: $\rho_{bd} = T(\rho_{bc}) = (|\psi_1^\uparrow\rangle \otimes |\psi_2^\uparrow\rangle)(\langle\psi_1^\uparrow| \otimes \langle\psi_2^\uparrow|)$ If now we attempt to trace ρ_{bd} over the subsystem associated with the edge e_d, we will obtain an incorrect result for ρ_b, namely $|\psi_1^\uparrow\rangle\langle\psi_1^\uparrow|$. The resolution is well known. Since a classical observer located on the edge e_b is not aware of the result of the intervention at the second vertex, for him the density matrix ρ_{bd} has evolved from ρ_{bc} by an operator $\tilde{T}$ which includes all possible outcomes of the measurement: $\tilde{\rho}_{bd} = \tilde{T}(\rho_{bc}) = \sum_{s=\uparrow,\downarrow} P_2^s \rho_{bc} P_2^s$. Tracing out the d-subsystem in the expression for $\tilde{\rho}_{bd}$, we obtain the correct expression for ρ_b, namely $\rho_b = 1/2(|\psi_1^\uparrow\rangle\langle\psi_1^\uparrow| + |\psi_1^\downarrow\rangle\langle\psi_1^\downarrow|)$.

Now we give another general prescription for computing the density matrix on an edge e_i from an arbitrary locative slice L containing this edge. We first compute a density matrix $\tilde{\rho}L$ for the slice L. This prescription has to deal with the possibility that some of the intervention operators used are from vertices that are not to the past of the edge e_i. As we saw in the example above, the density matrix at e_i cannot reflect the knowledge of the outcome of interactions at vertices that are not to the causal past of e_i.

This density matrix is computed from the initial data by applying intervention operators for the events in the past of L as before. But now, we will consider two types of events in the past of L: those that are to the past of e_i and those that are not. For the events that are to the past of the edge e_i, we use our regular intervention operators without a summation over the set of possible outcomes: $\rho \mapsto \sum_m A_{\mu m} \rho A_{\mu m}^\dagger$. We do not sum over the outcomes in this case precisely because the outcome is in fact known at e_i. For the events that are to the past of the slice L but not to the past of the edge e_i, we use operators that sum over all possible outcomes: $\rho \mapsto \sum_{\mu m} A_{\mu m} \rho A_{\mu m}^\dagger$. This time, of course, the summation is

there because the outcome cannot be known at e_i since these events are not to the past of e_i.

After we have obtained $\tilde{\rho}_L$, we trace out those subsystems associated with edges in L except for e_i to obtain the density matrix $\tilde{\rho}_i$. This is the density matrix associated with our preferred edge e_i, as computed from the slice L. The independence of the result on the choice of L is expressed in the following proposition:

Proposition 3.7. *Let e_i be an edge in the dag G. The density matrix ρ_i associated with the edge e_i does not depend on the choice of locative slice used to compute it, provided that we use the second prescription to compute it.*

Proof: We have already demonstrated that to any edge e_i, there is a unique least locative slice M_i containing e_i. Let ρ_i be the density matrix for the edge e_i as computed from the least locative slice and let $\tilde{\rho}_i$ be the density matrix for the same edge but computed from an arbitrary locative slice, say L, containing e_i. We will prove the proposition by showing that $\rho_i = \tilde{\rho}_i$.

First note that M_i being less than L implies that there is a set V of events between M_i and L. The plan is to remove the effect of these events and show that, at each stage, the density matrix is unaffected. We begin by picking a maximal event, say k, with the intervention operator T_k. Since k is maximal and hence acausal with all other maximal elements of V, as well as with all the maximal elements to the past of e_i, the intervention operator at k commutes with all the intervention operators at the vertices just mentioned. Thus, we can choose the intervention operator T_k to be the outermost, i.e. the density matrix ρ_L obtained by propagating to L can be written as

$$\rho_L = T_k(\rho')$$

where ρ' is the density matrix on the (locative) slice obtained by removing the edges to the future of k from L and adding the edges to the past of k. Using the explicit general form for an intervention operator,

$$\rho_L = \sum_{\mu,m} A^{(k)}_{\mu,m} \rho' A^{\dagger(k)}_{\mu,m}.$$

In order to obtain the density matrix $\tilde{\rho}_i$, we trace over all Hilbert spaces associated with edges in L except e_i. In particular, we trace over the outgoing edges associated with k. Now we can use the cyclic property of trace and rewrite the expression for $\tilde{\rho}_i$ as,

$$\tilde{\rho}_i = Tr\left(\sum_{\mu,m} A^{\dagger(k)}_{\mu,m} A^{(k)}_{\mu,m\rho'}\right).$$

Now we use the identity

$$\sum_{\mu m} A^{\dagger}_{\mu m} A_{\mu m} = I$$

to get

$$\tilde{\rho}_i = Tr(\rho')$$

We have eliminated the effect of the intervention operator at k. Proceeding inductively we can peel off the intervention operators associated with the rest of the vertices in V, thus

$$\tilde{\rho}_i = \rho_i .$$

□

A similar argument for the case of a simple system represented by the dag in Fig. 2 is contained in (Peres, 2000b).

4. CONCLUSIONS

We have presented an axiomatic system for the analysis of quantum evolution. The dynamics is local as to preserve causality, but at the same time entanglement of separated quantum systems is faithfully represented. Our work also suggests a natural extension of the notion of quantum history. Restricting the intervention operators at the vertices of our graph G to be projection operators we can consider G to denote a particular history within a set of histories. This relaxes the usual linear ordering of events considered in the literature thus far.

ACKNOWLEDGMENTS

The authors would like to thank NSERC for its financial support. The paper (Markopoulou, 2000), which led to our initial consideration of these ideas, was pointed out to us by Ioannis Raptis. We have benefited from useful discussions with Samson Abramsky, Fay Dowker, Chris Isham, and Fotini Markopoulou. We would like to especially thank Rafael Sorkin for a lengthy discussion on causal sets and related topics.

REFERENCES

Blencowe, M. (1991). The consistent histories interpretation of quantum fields in curved space-time. *Annals of Physics* **211**, 87–111.

Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. (1987). Spacetime as a causal set. *Physical Review Letters* **59**, 521–524.

Davies, E. B. (1976). *Quantum Theory of Open Systems*, Academic Press, New York.

Gell-Mann, M. and Hartle, J. B. (1993). Classical equations for quantum systems. *Physical Review D: Particles and Fields* **47**, 3345–3382.

Griffiths, R. B. (1996). Consistent histories and quantum reasoning. *Physical Review A* **54**, 2759–2774.

Markopoulou, F. (2000). Quantum causal histories. *Classical and Quantum Gravity* **17**, 2059–2077.

Markopoulou, F. and Smolin, L. (1997). Causal evolution of spin networks. *Nuclear Physics B* **508**, 409–430.

Nielsen, M. and Chuang, I. (2000). *Quantum Computation and Quantum Information*, Cambridge University Press, Cambridge, UK.

Omnès, R. (1994). *The Interpretation of Quantum Mechanics*, Princeton University Press, Princeton, NJ.

Peres, A. (1995). *Quantum Theory: Concepts and Methods*, Kluwer Academic, Norwell, MA.

Peres, A. (2000a). Classical interventions in quantum systems: I. The measuring process. *Physical Review A* **61**, 022116.

Peres, A. (2000b). Classical interventions in quantum systems: II. Relativistic invariance. *Physical Review A* **61**, 022117.

Plotkin, G. D. (1976). A powerdomain construction. *SIAM Journal of Computing* **5**(3), 452–487.

Penrose, R. and MacCallum, M. A. H. (1972). Twistor theory: An approach to the quantization of fields and spacetime. *Physics Reports* **6C**, 241–315.

Preskill, J. *Quantum Information and Computation. Lecture Notes*, California Institute of Technology, CA.

Raptis, I. (2000). Finitary spacetime sheaves. *International Journal of Theoretical Physics* **39**, 1703–1720.

Sorkin, R. (1991). Spacetime and causal sets. In *Relativity and Gravitation: Classical and Quantum*, J. D'Olivo *et al.*, ed., World Scientific, Singapore,