

A branching space-times view on quantum error correction

Thomas Müller*

25th May 2007

Abstract

In this paper we describe some first steps for bringing the framework of branching space-times to bear on quantum information theory. Our main application is quantum error correction. It is shown that branching space-times offers a new perspective on quantum error correction: as a supplement to the orthodox slogan, “fight entanglement with entanglement”, we offer the new slogan, “fight indeterminism with indeterminism”.

1 Introduction

Branching space-times (BST; Belnap 1992) is a general formal framework for describing indeterminism in a relativistic setting. We have good reason to assume that our world is both relativistic (the speed of light is the limiting velocity for signal transmission) and indeterministic (the way things are now does not always fix uniquely how things will be in the future). Branching space-times gives a formal model of how these could both be true.

Indeterminism means that certain situations allow for more than one possible continuation in their future. In a Newtonian world, this might be spelled out in terms of instantaneous states of the whole world across all of space. Relativity theory appears not to allow for an invariant notion of “instantaneous state of the world”. Thus, in the spirit of relativity theory, indeterminism must rather be a local affair. Accordingly, branching space-times pictures basic indeterministic happenings as *basic transitions* that consist of a single point-like event and one of its multiple possible causal futures (possible future light cones of the event).

*Institut für Philosophie, Lennéstr. 39, 53113 Bonn, Germany; e-mail: Thomas.Mueller@uni-bonn.de

More generally, in the framework of BST all indeterministic happenings can be viewed as indeterministic transitions that can be accounted for through the basic indeterministic transitions that are causally involved in them. Along these lines, the theory of transitions in BST has been used to give an analysis of causation (Belnap, 2005) and of single-case probabilities (Weiner and Belnap, 2006; Müller, 2005; Belnap, 2007).

Apart from such general applications, the formal framework of BST can also be used as a background for the formulation of specific physical theories. An obvious candidate is quantum mechanics, since that theory has both an indeterministic aspect, in the form of measurements and their outcomes, and is in need of a spatio-temporal setting, in view of the phenomenon of quantum correlations. A number of studies has shown the usefulness of a BST perspective on quantum mechanics (Belnap and Szabò, 1996; Placek, 2000; Müller and Placek, 2001; Müller, 2002). A recent result (Müller, 2007) furthermore shows how the histories approach to quantum mechanics (Griffiths, 1984; Gell-Mann and Hartle, 1990, 1993; Griffiths, 2003) can be interpreted within a subset of BST, viz., the branching time formalism. Thus there is some hope of arriving at a histories-type formulation of relativistic quantum mechanics based on BST.

One of the fastest growing areas of contemporary physics, with many promising potential applications, is the area of quantum computation and quantum information (cf. Nielsen and Chuang, 2000, for an overview). That field of quantum theory has not yet been approached from a BST perspective. Such an approach, however, appears promising: Firstly, from a general point of view, a quantum computer is an indeterministic spatio-temporal system, and describing such a computer in a formal framework unifying these two basic aspects may prove to be beneficial. Secondly, there are specific issues in quantum information theory in which indeterminism plays an important role and for which the formal tools of BST may provoke fresh insights. Two such fields are the theory of one-way quantum computation (Raussendorf and Briegel, 2001; Browne and Briegel, 2006) and the theory of quantum error correction (Nielsen and Chuang, 2000, Chap. 8–10).

In this paper, we will focus on the latter issue, i.e., quantum error correction, as a first step of bringing together the fields of branching space-times and quantum information. In order for the discussion to be accessible for readers familiar with only one of the sides, our presentation is perhaps more verbose than might be usual, and some readers may wish to skip a few paragraphs introducing material with which they are already familiar. The overall aim of this paper is to advertise methodological pluralism. We do not prove any new theorems or provide new error correcting codes. Nor do we claim that a BST view on quantum information is the only viable position. Rather, we wish to show that a BST perspective is *a* viable position, which may prove to be illuminating for conceptual issues of quantum information in general and quantum error correction in particular.

The paper is structured as follows: Section 2 reviews some key concepts of branching space-times. Section 3 then gives a brief overview of some aspects of quantum information theory and introduces the subject of quantum error correction. Finally, in Section 4 we present our branching space-times based view on quantum error correction.

2 Branching space-times

Branching space-times takes seriously the view that our world, the world we live in, is both indeterministic and relativistic. There have long been available formal frameworks for picturing our world separately as relativistic (the theory of Minkowski space-time, dating back to Minkowski 1908) and as indeterministic (the Prior-Thomason theory of branching time; cf., e.g., Thomason 1970). Formally, both these theories are based on partial orderings, i.e., transitive and anti-symmetric relations that are generally not total (linear) orderings, but have a more complicated structure.

Minkowski space-time is based on a partial ordering (the relation of causal accessibility) among point-like events. Not all of these events are connected by the ordering relation; events that are not thus connected are called *space-like related*. They coexist in a single space-time despite not being causally accessible one from another. No notion of indeterminism is forthcoming within Minkowski space-time.—Prior-Thomason branching time is based on a tree-like partial ordering among possible instantaneous states of the whole world. Maximal branches (maximal linearly ordered subsets of the tree-like ordering) represent possible courses of events; they are called *histories*. If two world-states are not connected by the ordering relation, this means that they belong to alternative, incompatible histories. A branching-time model can represent indeterminism simply by incorporating more than one history.

Belnap's (1992) theory of BST unifies these two formal models by blending together the partial orderings on which they are based. In BST, there is a single partial ordering that can be interpreted as the relation of *possible causal precedence* among possible point events. That ordering is presumed to be continuous and unbounded, thus capturing the key features of Minkowski space-time, and it also allows for a notion of indeterminism. According to BST as well as according to branching time, our world contains entities that cannot occur together in one course of events, or history: that is how our world must be like if there is indeterminism. In branching time, a history (possible total course of events) corresponds to a maximal linear path through the given tree-like partial ordering. In BST, a history cannot be linearly ordered, since that would rule out space-like related events. Thus, histories in BST are defined in a way that allows each history in our

world to be, e.g., a full Minkowski space-time.

Local, relativity-friendly indeterminism means that histories must split locally, at individual point events. Above such an indeterministic *splitting point*, there are at least two possible future light cones. Figure 1 pictures a situation involving two binary (+/-) splitting points arranged in a space-like manner, leading to four histories. A splitting point together with one of its possible futures represents a single *basic transition* in the BST model.




Figure 1: Representation of a BST partial ordering with two binary splitting points (outcomes +/-) arranged in a space-like manner, leading to four histories. Points labelled identically occur in all the histories in which that label occurs.

To give a flavour of the machinery of BST, we give some of the key definitions. Technically, BST starts with two primitive notions: “Our world”, W , whose members are defined as *point events*, and \leq , the “causal order” on W . It is assumed that \leq is a dense partial order on W with no maximal elements. A *history* h is a maximal directed set, where a set is *directed* iff it contains an upper bound for each pair of its members. Point events are *space-like-related* iff they are distinct, not causally ordered and share a history. H denotes a set of histories, and $H_{(e)}$ denotes the set of all histories in which the point event e occurs.

An *outcome chain* O is a nonempty and lower bounded chain, where a set is a *chain* if each two of its members are comparable by \leq . One can prove that $O \subseteq h$ for some history h . It is assumed that O has always a unique infimum $\inf O$, and it is provable that given $e \in h$, there is an O such that $O \cap h \neq \emptyset$ and $e < O$ and $e = \inf O$. It is also assumed that every upper bounded chain has a supremum in every history to which it belongs. An *initial event* I is a set of point events all of which are members of some one history. By the *prior choice postulate*, for every

$O \subseteq h_1 - h_2$ there is a point $e \leq O$ that is maximal in the intersection $h_1 \cap h_2$. These assumptions fix the topology of models of BST.

Two histories are *undivided at* e , $h_1 \equiv_e h_2$, iff e belongs to their intersection, but is not maximal therein. h_1 is *separated* from h_2 at e , written $h_1 \perp_e h_2$, iff e is maximal in $h_1 \cap h_2$; such an e is a *splitting point* for histories h_1 and h_2 . Note that histories can split at more than one point, like the upper left and lower right histories in Figure 1. This feature makes it possible to define notions of modal and probabilistic correlation that can be used, e.g., for applications in quantum mechanics (Belnap, 2003; Müller, 2005).

It follows from the prior choice postulate that undividedness-at- e is an equivalence relation. The corresponding partition of $H_{(e)}$ is denoted Π_e . Given $e \in h$, $\Pi_e(h)$ is that member of Π_e to which h belongs. A *transition* is a pair $\langle I, O \rangle$ of an initial and an outcome, written $I \rightsquigarrow O$, such that I entirely precedes O . *Basic transitions* $t = e \rightsquigarrow H$ are from a point e to one of the elements $H \in \Pi_e$ of the partition of $H_{(e)}$. Within the continuous ordering of a BST model, these basic transitions indicate the concrete point events at which indeterminism manifests itself. They are therefore the basic building blocks both for an analysis of causation in BST (cf. Belnap, 2005) and for a BST-based probability theory (cf., e.g., Müller, 2005).—For further technical details of BST we refer to Belnap (1992, 2005).

3 Quantum information

Quantum theory has been, and still is, the source of many mysteries. Quantum information theory offers an interesting view on the properties of quantum systems since it treats quantum phenomena such as entangled states not as conceptual mysteries to be solved, but as physical resources to be exploited. While the practical design of actual, working quantum computers is still an area of frontier research, much progress has been made with respect to the theoretical foundations of quantum information and quantum computation.¹ We believe that branching spacetimes can offer an alternative and potentially useful view on some of the concepts used in these disciplines.

3.1 Classical and quantum information

Classical information In classical information theory, the basic unit of information and computation is the binary digit, or bit. A bit can be in one of two states, which are commonly denoted 0 and 1 and which correspond to an answer

¹A standard reference in the field is Nielsen and Chuang (2000). For a good, more recent overview of the area, cf. Eisert and Wolf (2006).

to a single yes/no question. Reading out the state of a bit provides at most an epistemic, but no ontological source of indeterminism: if the state of the bit is not known beforehand, the result of the read-out will be one of two epistemic alternatives. However, in the absence of errors or noise (cf. Section 3.3 below), the value that is read out is exactly determined by the state that the bit is in, so there is no indeterminism proper. In a branching space-times picture, reading out a classical bit therefore does not correspond to a splitting of alternatives. The state of a bit can be copied easily without affecting the state of the bit that is copied. Bits can be realised physically in many different ways, e.g., magnetically or via integrated circuits of various sorts, and technology is advancing fast, making for ever smaller and faster devices for storing classical information.

Quantum information In quantum information, the basic unit of information and computation is the so-called quantum bit, or qubit. Physical realisation of qubits is at present a subject of frontier research. A qubit corresponds to the state space of a two-state quantum system, which space is a two-dimensional Hilbert space. One set of orthogonal normalised basis vectors for that space, the so-called computational basis, is commonly denoted $\{|0\rangle, |1\rangle\}$, but there are infinitely many other possible bases. A general (pure) state $|\psi\rangle$ of a qubit can be written as a superposition of basis vectors:

$$|\psi\rangle = a|0\rangle + b|1\rangle, \quad (1)$$

where a and b are complex numbers such that $|a|^2 + |b|^2 = 1$. Thus, there is much more to the state of a qubit than an answer to a single yes/no question, even though a qubit still carries only a single bit of information in the sense of Shannon information theory (Shannon and Weaver, 1949).² At any rate, $|\psi\rangle$ is determined by two real-valued parameters, the modulus of a and a relative phase. Accordingly, there are many more operations that can be performed on qubits than on classical bits. In quantum algorithms, both unitary transformations and projective measurements are employed, and both classes of operations contain continuously many alternatives.

Reading out the state of a qubit poses problems of a fundamental nature: in general, it is impossible to determine the coefficients a and b with respect to any basis. What one can do effectively amounts to a measurement of the state with respect to one single basis (chosen by the experimenter), which process is commonly interpreted as projecting the qubit into one of the two basis states (e.g., $|0\rangle$ or $|1\rangle$), destroying the antecedent state $|\psi\rangle$. That way, the qubit in effect reveals no more than an answer to a single yes/no question.

²For a criticism of a different interpretation of the phenomenon of “dense coding”, cf., e.g., Griffiths (2002).

Furthermore, that answer is of a stochastic nature: except for the case in which the qubit is already in one of the two basis states, reading out a qubit is a truly indeterministic process. Even if the state is known perfectly (e.g., because the qubit was prepared in that very state), the result of reading out is not determined before it actually occurs. Thus, in a branching space-times picture, reading out a qubit in state $|\psi\rangle$ corresponds to a twofold splitting of alternatives, with probabilities given by the inner product of the state of the qubit and the respective basis vectors, which probabilities are equal to $|a|^2$ and $|b|^2$ in the $\{|0\rangle, |1\rangle\}$ basis. The outcome of such an ontological splitting due to measurement is normally known: e.g., after a spin measurement one knows the particle's spin; the notion of measurement is conceptually tied to the notion of a (known) measurement result. An ontological splitting the outcome of which is not known, would normally count as an error, as a disturbance of the system, not as a measurement (cf. Section 3.3 on the concept of an error). This epistemic aspect of the measurement is not part of the basic BST description, but it can be superadded to the BST structure, e.g., via some sort of labelling.

One might hope that the true state of a qubit could be determined more fully by making copies of the state $|\psi\rangle$ and subjecting these copies to measurement with respect to various, cleverly determined bases, but that route is blocked: for reasons of principle, the state of a qubit cannot be copied (“no cloning theorem”; Dieks 1982).

3.2 Classical and quantum operations

Classical information processing Classical information can be processed in a variety of ways. The best-known way proceeds in the manner of Boolean algebra. In any scheme, the state of a bit can be used many times, as it may be copied freely. Some of the basic Boolean operations are negation (turning the state of a single bit from 0 to 1 and vice versa) and logical “and” and logical “or”, both producing one output bit from two input bits in an irreversible manner. The result of these operations is perfectly determined by the input, so there is no indeterminism (no BST splitting) in classical Boolean computation. There are also other, well studied ways of employing classical information: classical computation can be done in a reversible manner (much like in quantum computation), and in classical random computation, indeterministic BST splitting plays an important role.

Quantum information processing Qubits allow for many more “logical” operations than Boolean algebra. In fact, any unitary operator can be viewed as a means of computation. This fact allows for specific quantum algorithms that are believed by many (but have not been proved) vastly to outperform any classical

analogue. Most famously, there is a fast quantum algorithm for factoring natural numbers (Shor, 1997). The existence of this and similar algorithms provides one of the strongest incentives to make quantum computation work in practice.

From a BST perspective, quantum computation via unitary operations is as unspectacular as classical Boolean computation: no indeterminism is introduced by performing unitary operations on quantum systems, so there is no BST splitting involved. In a certain sense, quantum computation is even *less* spectacular than classical Boolean computation: as the states of the qubits are only subject to unitary transformations, quantum computation is a completely reversible process. This is not normally so for classical computation (e.g., one cannot normally reconstruct the input of an “and” gate from its output), which fact points to a further, thermodynamic advantage of quantum computation: quantum computers do not heat up as much as classical computers.

It should be noted, however, that while the paradigm of reversible, unitary computation described in the previous paragraphs represents the mainstream view on quantum computation, there has been proposed a scheme for quantum computation operating on so-called graph states (Raussendorf and Briegel, 2001) that, at least according to most interpretations, makes heavy use of indeterministic splitting. That scheme of implementing quantum computation may in fact be one of the technologically most accessible options (Browne and Briegel, 2006; Gross and Eisert, 2006). In what follows, we will not comment on this development, which would merit a separate branching space-times analysis.

3.3 Errors and error correction

Errors Stored or processed information can be subject to corruption due to errors (also called “noise”), for various reasons. The simplest, well established model for describing such processes, already employed in the groundbreaking work on information theory by Shannon and Weaver (1949), is that of a “noisy channel”. The channel has an input side, at which known information enters, and an output side, at which the information is received, perhaps in corrupted form. A classical channel is specified by giving probabilities for various forms of corruption of information in transmission. One may think of such a channel as a physical wire carrying electrical signals, or as the mere passage of time, or as some more elaborate spatio-temporal arrangement. In a branching space-times picture, a noisy classical channel corresponds to indeterministic splitting, with the further proviso (not itself part of the basic BST description) that the outcome of that splitting be *unknown*. This epistemic fact is part of the concept of an error: known errors are not errors, but transformations that can usually be accounted for easily. E.g., if a channel has a certain probability of flipping a bit that is fed into its input side, we must assume that it is unknown at the output side whether the

bit was in effect flipped or not—otherwise, we could simply adjust for the flip if it occurred, and have a perfect, noise-free channel.

In classical computation, the two states of a bit may be realised physically in such a way that some means of error correction are quasi-automatic. E.g., if the state of a bit is encoded by a certain voltage (as in a digital computer), then choosing the voltage difference between the codes for 0 and 1 large enough will allow one to tolerate a fair amount of noise simply by counting a voltage at the input of a classical gate above a certain threshold as 1 and the rest as 0. If the gate emits the standard voltage at its output, this means that noise does not accumulate.

While such considerations may help to improve a device’s error characteristics, in many situations, the choice of a particular channel with particular error characteristics is forced upon one, so that any application must make do with given imperfections. Any digital storage medium or digital signal transmission scheme needs to take account of the possibility of random errors. Hard drives, DVDs, and LAN links all need to be (and of course, are) provided with means of recovering from errors. The same holds, even to a greater extent, for purported physical implementations of quantum computers, in which a broader approach to the notion of an error will be needed. In fact, just as the state space of a qubit allows for continuously many states, so there are continuously many possibilities of getting things wrong. If there were no means of error correction in quantum computation, it would be infeasible even to try to construct a working quantum computer.—We will discuss a general approach to quantum channels in Section 4 below.

Error correction In terms of the paradigm of the “noisy channel”, the idea of error correction is to introduce redundancies into the information that is fed into the input side of the channel, so that at the output side, these redundancies can help to arrive at a safe guess as to what the original data was, even if the received information differs from what was sent—which fact is, of course, unknown. Consider, e.g., the international radiotelephony spelling alphabet, according to which one indicates the letter “C” by the word “Charlie”: one transmits much more information than would be necessary, but even in a noisy environment, one can decode what was said with high accuracy.—Depending on the application, finding good error correcting codes can be something of an art, quite apart from tackling the mathematical problems involved. The book by Welsh (1988) provides a valuable overview of the classical theory.

Majority voting Perhaps the simplest error correcting code is to employ so-called majority voting. For classical information, a bit can either be transmitted correctly, or be flipped (turning 1 to 0 and vice versa). We consider a noisy channel

that has a probability p of flipping any one bit and no correlations among consecutive flippings. For this channel, transmitting a number of copies of each input bit provides an effective means of error correction: if one transmits each bit, e.g., three times through a channel with bit flip probability p , taking the majority vote of bits read out at the output side improves reliability in any case in which $p < 1/2$: the probability of getting things wrong is the probability of two or more bit flips, which is equal to $p^3 + 3p^2(1 - p)$. This can lead to a marked improvement—e.g., for $p = 0.1$ (one bit flip error in ten transmissions), the probability of getting the corrected output wrong is 0.028.

Majority voting is of course not a very good code for technical applications involving classical information—there are much more elaborate schemes that make better use of redundancies. However, majority voting already gives the flavour of what is involved in error correction. Furthermore, it is a scheme that is actually used in oral communication (e.g., repeating announcements multiple times at a noisy airport terminal), and it provides the basis for the simplest quantum error correction code.

Quantum error correction must appear to be much more difficult than error correction for classical bits. A straightforward analogue of the classical techniques is not possible. After all, e.g., the majority voting scheme relies on the classical possibility of copying the state of a given bit at the input side—but that is impossible for qubits, according to the no cloning theorem. Still, quantum error correction is possible by using quantum effects in a clever way.

4 Quantum error correction: a branching space-times view

Employing BST as a framework for describing quantum effects means to place emphasis on the indeterministic aspects of quantum theory. Like the “consistent histories” approach, a BST view treats the evolution of a quantum system as a stochastic process that generally allows for many different descriptions. That view deviates from the mainstream by taking a specific stance with respect to the interpretation of quantum theory. The main aim of the sketch we give here is to show that BST offers a viable picture, thus broadening the range of available points of view.

It is well known that quantum mechanics is troubled by deep problems of interpretation. If quantum mechanics is true, then the Schrödinger equation should govern the evolution of the universe. The Schrödinger equation describes a unitary, deterministic evolution. However, some of the most striking quantum phenomena appear to be fundamentally indeterministic. In order to account for this

fact, quantum theory is supplemented with a probabilistic algorithm describing the outcome of so-called measurements. The interrelation between these two parts of quantum theory is still far from clear. There is an immense discussion about the issues involved, which we leave behind (cf., e.g., Bub, 2000). One aspect of the so-called measurement problem is the question of the criteria for a “measurement”, i.e., for something that should appropriately be modelled via the indeterministic aspect of quantum mechanics and not deterministically. If quantum theory really is a fundamental theory, it appears to be possible to push the boundaries ever further out and to describe what appears as measurement from one perspective, as a deterministic evolution of measured system plus measurement apparatus. There seems to be no non-arbitrary way of singling out the “real” measurements, even though in some cases, there is widespread agreement among the experts.

The ontology of quantum effects thus appears to be a shifting business. Pragmatically it is often useful to link ontology with epistemology at this point, i.e., to treat *unknown* influences on a system in the unitary “system plus environment” picture, while reserving the probabilistic algorithm for predictions of influences whose result will be *known*. This attitude is widespread among the quantum information community. In measurements, to quote a metaphor suggested by one of the referees for this paper, “you force the system to decide” one way or another. On this view, the system really evolves indeterministically, but the result of the evolution will be known. The information resulting from such measurements is often employed as classical information in an overall computational scheme.—On the other hand, the evolution of a quantum system subject to random disturbances, the results of which will not be known, is standardly modelled from an “open system” point of view, in which one focuses on a unitary evolution of system plus environment, through which the system becomes entangled with the environment. Any air of stochasticity in the state of the system alone comes about by tracing out the environmental degrees of freedom. On this view, the system plus its environment evolve deterministically, and the indeterministic aspect of the situation is treated as having a purely epistemic source.

This point of view has obvious merits, and it is strongly connected to the general “open systems” point of view of quantum optics, which supplies many of the impulses for quantum information theory. We have no fundamental qualms with this epistemic-ontic link, but by presenting our BST view we wish to point to an alternative conception that denies the mentioned interpretational link.

Branching space-times places heavy emphasis on indeterminism. The formal framework itself admits both an epistemic and an ontic interpretation of the splittings represented in BST models, even though much of the initial motivation for the development of BST comes from an ontic interpretation (cf. Belnap, 1992, 2005).

If a system evolves deterministically, its evolution will be described as a single

history, and a BST description will not be illuminating. With respect to quantum information and quantum error correction, it is thus important to determine which indeterministic happenings should be modelled. As was argued in the previous paragraphs, identifying indeterminism in a quantum system may to some extent be a question of preference rather than a matter of fact, and at least the quantum histories community shows that a number of people prefer an emphasis on a stochastic evolution. The BST view on quantum error correction deviates from the standard view exactly by identifying more indeterministic happenings than is usual. As was stated in the introduction, this is not meant to rule out the standard view, but rather to supplement it, in the spirit of methodological pluralism. Our BST picture mainly deviates from the standard view in the way we approach the concept of an error, as will be pointed out in Sect. 4.2 below.

In classical information theory, a unit of information is in one of a discrete set of alternative states—for a bit, these are just the two possible states, 0 and 1. An error, or “getting things wrong” with respect to classical information means an unknown mixing up these states. For a bit, there is only one type of error, which is an (unknown) bit flip; more generally, one can describe a classical channel in terms of the probabilities of the various possible errors. For quantum information, the situation is different, since there is a continuum of states that a quantum system such as a qubit may be in. Errors can effect the state more or less drastically, while quantum algorithms may rely crucially on the exact state of a qubit. Error correction is thus both more important and more difficult in the quantum domain.

In this Section we will describe a general approach to “what may happen to a quantum system” in terms of quantum channels, and for one specific channel we will sketch a quantum error correction scheme that shows some similarities to majority voting as discussed above. For both topics, we will first give a more or less standard presentation as is customary in a quantum information setting and then give our branching space-times view. Thus, Section 4.1 introduces quantum channels, while Section 4.2 gives a BST view on quantum channels; Section 4.3 discusses a simple quantum error correction scheme for the so-called bit-flip channel, and Section 4.4 gives a BST view on that error correction scheme.

4.1 Quantum channels

A quantum channel is a general operation on the state of a quantum system. Such an operation can be described mathematically as a completely positive (CP), trace-

preserving (TP) map³

$$\Phi : \rho \mapsto \rho' = \Phi(\rho). \quad (2)$$

Unitary transformations describing the deterministic evolution of the system itself are but a very special case. A more general class of channels is the class of unital channels. Many interesting quantum channels arising in the study of quantum errors are unital (or bistochastic), meaning that they map the state I on I ,

$$\Phi(I) = I. \quad (3)$$

A unital channel for a single qubit may be pictured in analogy with a classical channel, viz., as a convex combination of various unitary transformations of the system (Birkhoff's theorem). This picture breaks down for channels of higher dimensions as well as for non-unital channels.⁴

In the general case, there are at least two useful ways of picturing the action of a quantum channel. One is to think of the (generally, not unitary) transformation of the system affected by the channel as the partial trace over the result of a unitary transformation of the system plus some environment. It turns out that for a system of dimension d , an environment of dimension d^2 is enough to model all quantum channels, and that the environment may be chosen to be in a standard pure state (Stinespring dilation). This system plus environment picture has great appeal since it gives a non-stochastic (unitary) reading of unknown disturbances of a system's state in terms of decoherence or entanglement of the system with its environment, thus preserving the ontic-epistemic link mentioned above.

The other useful picture of a quantum channel focuses on the system only. By the Kraus representation theorem, a completely positive, trace preserving map Φ can be described in terms of a (non-unique) set of operators E_k such that

$$\rho' = \sum_k E_k \rho E_k^\dagger, \quad (4)$$

where

$$\sum_k E_k^\dagger E_k = I. \quad (5)$$

(Bistochasticity corresponds to the additional condition that $\sum_k E_k E_k^\dagger = I$.) In terms of density matrices

$$\rho_k = E_k \rho E_k^\dagger / \text{Tr}(E_k \rho E_k^\dagger) \quad (6)$$

³Preservation of the trace may be assumed without loss of generality—scenarios in which the trace is not preserved correspond to descriptions of parts of a system for which the trace *is* preserved.

⁴Cf. Landau and Streater (1993).—Useful literature on quantum channels includes Nielsen and Chuang (2000, Chap. 8), Griffiths (2002), and Eisert and Wolf (2006).

and probabilities

$$p_k = \text{Tr}(E_k \rho E_k^\dagger) \quad (7)$$

the Kraus representation may be written as

$$\rho' = \sum_k p_k \rho_k, \quad (8)$$

which suggests interpreting the action of the channel as replacing the system's initial state ρ by the state ρ_k with probability p_k .

4.2 Quantum channels: a BST view

The BST view on a quantum channel takes at face value the mentioned interpretation of the Kraus representation of a given quantum channel: The action of a channel with representation $\{E_1, \dots, E_n\}$ is pictured as an n -fold indeterministic splitting at the spatio-temporal location of the channel output, corresponding to replacement of the input state ρ by the various ρ_k with respective probabilities p_k . As the channel is used to model an unknown influence on the system, the splitting in the BST picture needs to be unknown, which fact may be indicated by the absence of identifying labels for the outcomes.⁵

The picture suggested here is very simple, and may appear to be open to two obvious challenges. It may be well to address these here in order to show that our proposal is in fact tenable.

Non-uniqueness The Kraus representation of a given CP map is not unique. Thus, what could be the status of a given BST description of a quantum channel?

This type of problem (if it is a problem) pertains to any histories-type approach to quantum mechanics. The standard statistical picture gets around this difficulty since if one singles out measurements explicitly, one thereby also singles out a unique basis for the description of the system: the measurement device provides a basis (the “pointer basis”) non-arbitrarily in terms of the possible measurement results (“pointer states”). However, that non-arbitrariness comes at the high conceptual cost of the measurement problem, as pointed out above. Any approach to quantum theory that includes stochasticity without measurement is confronted with the problem of non-uniqueness.

⁵We stress again the fact that the basic machinery of BST as sketched in Section 2 does not include probabilities nor a notion of “epistemic state”. Probabilities may be added as in Müller (2005), which approach employs sets of basic transitions with the same initial as the basis for introducing probability spaces. A good technical implementation of the epistemic dimension is still a desideratum. For a first step in that direction cf. the notion of K -extension in Placek (2002).

With regard to this problem, we side with other histories-type approaches to quantum mechanics: The various possible stochastic descriptions of a system must be taken to have the same ontological status, but they may differ with respect to their epistemic merits or practical usefulness.

Consistency of the approach Other histories-type approaches (Gell-Mann and Hartle, 1993; Griffiths, 2003) stress the crucial role of a consistency condition, which amounts to using stochastic descriptions of a system’s dynamics that lead to orthogonal alternatives. Thus, for each time step of the evolution pictured in a consistent histories framework, the total system considered has to contain a subsystem that effectively keeps a record of what happened at that time step. A consistent histories approach to quantum channels (Griffiths, 2002, 2005) therefore forces one to include the environment in the description of the dynamics, making our approach appear questionable. Furthermore, while the Kraus operators, in view of (5), have some similarity with projective decompositions of the identity as employed in the consistent histories approach, they do not generally lead to orthogonal alternatives. Is the BST picture of quantum channels consistent?

Our approach differs from other histories-type approaches in focussing on the system under consideration exclusively. As discussed in Müller (2007), in a branching framework the consistency condition is not a precondition for assigning probabilities to alternatives, but rather a precondition of a type of coarse-graining that may be problematic in itself. If one focuses on a system without its environment, one does not describe records. One may move to a “system plus environment” description if an application demands that records be included in a description, and in that case, the consistency condition is crucial. For our purposes, it is sufficient to note that any quantum channel may be described in the “system plus environment” picture if one wishes to do so, as discussed in Section 4.1.

4.3 Quantum error correction

We pointed out above that quantum error correction poses a number of technical and conceptual challenges. Since the mid-1990s, a number of increasingly sophisticated quantum error correction schemes have been proposed. As in classical error correction, the basic idea is to introduce redundancies via coding and exploit these redundancies via decoding after noise may have affected the system. It turns out that more redundancy is needed for protecting qubits than for protecting classical bits—in order to protect a qubit against arbitrary errors, a five-qubit encoding has to be employed at minimum, while the simple majority voting scheme sketched above employed three bits per classical bit. The idea of *resilient computation* is that it is possible to do quantum operations on encoded states, thus

minimising the amount of coding and decoding necessary. The book by Nielsen and Chuang (2000, Chap. 8–10) gives an overview of the relevant techniques and results and also describes the mentioned five-qubit “stabilizer” code; cf. also Eisert and Wolf (2006).

We will now describe the key component of Shor’s (1995) error correction scheme, which is conceptually the simplest option. Overall, the code uses nine qubits per qubit to protect against arbitrary errors on a single qubit. The code consists of two concatenated layers. The first layer protects the initial qubit against so-called bit-flip errors by a three-qubit encoding, as already suggested by Peres (1985). A bit flip is a disturbance of the system that turns the state $|0\rangle$ into $|1\rangle$ and vice versa, and that affects a general state $|\psi\rangle$ via

$$|\psi\rangle = a|0\rangle + b|1\rangle \quad \mapsto \quad a|1\rangle + b|0\rangle. \quad (9)$$

While for a single classical bit, flipping was the only possible disturbance of a bit’s state, the state of a qubit can be altered in continuously many more ways. The Shor code accounts for these other types of disturbance via a second layer of encoding. That second layer first transforms the three output qubits of the first layer and then encodes each of them once more in the same manner, leading to a nine-qubit code. Even though the code is thus put together in a way that appears to offer protection against two discrete types of errors only, it actually offers protection against arbitrary one-bit errors, as described, e.g., in Nielsen and Chuang (2000, Chap. 10). From a conceptual point of view, the bit flip component already contains the main ingredients of a full error correcting code such as Shor’s.

We will thus describe error correction for the bit-flip channel, which is an especially simple unital channel that may be pictured in analogy with a classical noisy channel. Operating on input

$$|\psi\rangle = a|0\rangle + b|1\rangle, \quad (10)$$

the bit-flip channel yields as output either the bit-flipped state, $a|1\rangle + b|0\rangle$, with (hopefully small) probability p , or the correct, original state $|\psi\rangle$ (with probability $1 - p$). In the BST picture, the bit-flip channel corresponds to a twofold indeterministic splitting, while the standard “system plus environment” picture can make do with a one-qubit environment.

How can one correct for the kind of errors that are introduced through a bit-flip channel? Implementing a straightforward majority voting scheme is out of the question in view of the no cloning theorem. The remedy is to use $|\psi\rangle$ to create, not three *copies*, but an *entangled* three-particle state: from $|\psi\rangle$ it is possible, using two additional qubits in the $|0\rangle$ state and unitary transformations (two CNOT gates), to construct the encoding state

$$|\psi_c\rangle = a|000\rangle + b|111\rangle. \quad (11)$$

The individual particles in that entangled state can then be sent through the bit-flip channel independently. At the output end of the channel, measurement of any single particle's state would destroy the entanglement. However, one can use multi-particle measurements to determine a so-called error syndrome, which gives information on which error, if any, has occurred. The state space of the three entangled qubits used for coding has dimension 8. One two-dimensional subspace is spanned by the basis vectors $|000\rangle$ and $|111\rangle$. That subspace corresponds to the original, undisturbed coding or the—hopefully rare (probability p^3)—occurrence of three separate bit flips on all three particles. Similarly, the subspace spanned by the basis vectors $|100\rangle$ and $|011\rangle$ corresponds to a bit flip on the first particle or two separate bit flips on particles 2 and 3; analogously for the spaces spanned by $|010\rangle$ and $|101\rangle$, and by $|001\rangle$ and $|110\rangle$, respectively.

By a measurement that projects the three particle output state into one of these four two-dimensional subspaces (employing the projectors $|000\rangle\langle 000| + |111\rangle\langle 111|$ and so forth), one does not destroy the superposition, but through the measurement result one determines which—if any—single-bit error occurred. This error syndrome information can then be used to correct for the error, through a unitary transformation on the qubit that was flipped.

As with majority voting, this scheme can fail. It will fail if two or more bits were flipped, since then the error syndrome will suggest a different kind of disturbance. However, in all cases in which $p < 1/2$, the error correction scheme improves reliability, with the probability of getting things wrong again given by $p^3 + 3p^2(1 - p)$.

Preskill, one of the pioneers of quantum information theory, illustrates the idea behind this error correction scheme by the slogan, “we can *fight entanglement with entanglement*” (Preskill, 1998, 386). The picture is that bit flips or other noise can be understood from the standard “system plus environment” point of view as deterministic processes of entanglement between the quantum computer's qubits and their environment. The input coding then employs *additional* entanglement as a resource that helps one to cope with the unwanted entanglement due to noise: the entangled three-particle state used for coding the input allows for (indeterministic) error syndrome measurement and thus, for error correction, as described.

4.4 Quantum error correction from a BST perspective

Assuming the standard ontic-epistemic link mentioned above, a quantum error *must* be described as the unitary evolution of the system plus some environment, not as a case of indeterminism. This picture is of course consistent, but there is some tension with the intuitive notion of an error, which suggests something that “just happens”, pointing to indeterminism. Preskill's slogan, which is descriptive of the standard view, points to the fact that within the quantum information

community, indeterminism is recognised only for experimenter-induced measurements. Thus, e.g., the syndrome measurement described in the previous Section is a well-recognised type of indeterminism, but errors are not.

Branching space-times offers a different, complementary picture by focussing not on the notion of entanglement of system and environment, but on the notion of indeterminism. From a BST perspective, errors *are* indeterministic disturbances of the system's state: according to the BST point of view, the bit-flip channel acts indeterministically as laid out in Section 4.2, and an epistemic interpretation of the respective indeterministic splittings is not forced upon us. Thus in the following, we treat this indeterminism as ontic, like in the case of measurements. For the coding described above, in an indeterministic picture of the system's dynamics there are three indeterministic two-way splittings (bit flip: yes or no), which lead to a total of eight histories from the one history the input (given state ρ) started with.

The error syndrome measurement is again an indeterministic process that leads to a further four-fold splitting of histories, resulting in a total of 32 combinatorially possible histories. Now what about error correction? How can this blowing up of possibilities improve the situation? The key to the solution is that the knowledge of the outcome of the error syndrome measurement together with knowledge of the system's dynamics allows one partially to recover from the uncertainty about which history one is in that is introduced by the noisy channel. The eight-fold splitting due to the bit flip channel operating on the three qubits used for coding is epistemically inaccessible, as noted above: the concept of an *error*, or of noise, implies that one does *not know* whether a bit flip occurred or not. The error syndrome measurement, while ontologically a splitting of alternatives as well, has different epistemological consequences in view of its status as a *measurement*: the result of the splitting is *known*. Thus, one gains information even though there is a splitting of histories that one cannot control. Furthermore, the quantum dynamics of error syndrome measurement means that the information that one gains is usefully correlated with the unknown result of the previous splitting due to noise. While there are 32 ontologically possible histories in the given set-up (eight-fold splitting due to the bit flip channel, followed by a four-fold splitting in each of these histories corresponding to the four possible results of the error syndrome measurement), only eight of these histories are dynamically possible, i.e., have a non-zero probability as given by the probabilistic algorithm of quantum theory. E.g., if no bit flip occurred, the syndrome measurement will definitely show the output three-particle state to lie in the subspace spanned by the basis vectors $|000\rangle$ and $|111\rangle$. Figure 2 illustrates how a single history in the error correction scheme for the bit-flip channel looks like from a BST perspective. The epistemic dimension of the set-up is indicated by the question marks after the (unknown) splitting due to the transmission of the qubits through the bit flip channel (flip?), in con-

trast to the determined outcome (syn 0–3) after the (later) splitting due to the error syndrome measurement.

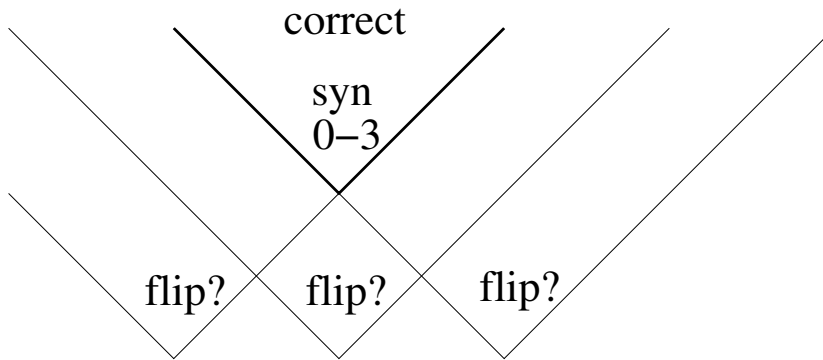


Figure 2: Generic illustration of one of the 32 ontologically possible histories involved in the three particle error correction scheme for the bit-flip channel.

Figure 3 illustrates the effect of the quantum dynamics on the probabilities of histories resulting from the syndrome measurement. Only one of the four ontologically possible histories is not crossed out, indicating non-zero probability. Correcting for the error indicated by the syndrome in this case would effectively restore the input state $|\psi_c\rangle$ used for coding.




Figure 3: Illustration of the quantum probabilities for results of error syndrome measurement, given that in fact, a bit flip on the second particle occurred. Crossed-out histories have zero probability.

The result of the error syndrome measurement does not completely reveal which of the eight dynamically possible histories one is in: if all three bits are flipped, the result of the syndrome measurement will be the same as if no bit

had been flipped. Still, there is a net gain: altogether, the syndrome measurement improves one's epistemic status from being in one of eight possible histories without any knowledge of which history one is in, to being in one of eight dynamically possible histories with two classical bits of information as to which one it is one is in. Thus additional indeterminism introduced through the syndrome measurement effectively improves one's epistemic status. Of course, the presence of noise means that one cannot recover the initially given state with perfect certainty. Practically, however, the probabilities of getting things wrong in the end are what matters, and these can improve drastically.

5 Summary

Quantum error correction has been known for a number of years. The conceptual basis of quantum error correction is usually seen in the possibility of employing entanglement as a resource that helps one to counter unwanted disturbances. Branching space-times offers a fresh and possibly helpful, different conceptual picture of what is involved in quantum error correction. By focussing on the indeterminism present, our branching space-times analysis of quantum error correction suggests to supplement Preskill's slogan, "fight entanglement with entanglement", with a new slogan: from a branching space-times perspective, the key to quantum error correction is to "fight indeterminism with indeterminism".

Acknowledgements

We would like to thank the European Science foundation and the Jagiellonian University, Kraków, for kind support for the workshop on Branching Space-Times (October 24–25, 2005), at which a previous version of this paper was presented. Many thanks to the workshop audience for helpful comments, and to two anonymous referees for their detailed critical remarks on previous versions of the paper. Special thanks to Nuel Belnap, Jens Eisert, Tomasz Placek, and Miklós Rédei.

References

- Belnap, N. (1992). Branching space-time. *Synthese*, 92:385–434.
- Belnap, N. (2003). No-common-cause EPR-like funny business in branching space-times. *Philosophical Studies*, 114:199–221.

- Belnap, N. (2005). A theory of causation: Causae causantes (originating causes) as inus conditions in branching space-times. *British Journal for the Philosophy of Science*, 56:221–253.
- Belnap, N. (2007). Can propensities be single-case causal probabilities? *Studies in the History and Philosophy of Modern Physics*, this issue.
- Belnap, N. and Szabò, L. E. (1996). Branching space-time analysis of the GHZ theorem. *Foundations of Physics*, 26:989–1002.
- Browne, D. E. and Briegel, H. J. (2006). One-way quantum computation — a tutorial introduction. quant-ph/0603226.
- Bub, J. (2000). *Interpreting the Quantum World*. Cambridge: Cambridge University Press.
- Dieks, D. (1982). Communication by EPR devices. *Physics Letters A*, 92(6):271–272.
- Eisert, J. and Wolf, M. M. (2006). Quantum computing. In Zomaya, A. Y., editor, *Handbook of Nature-Inspired and Innovative Computing: Integrating Classical Models with Emerging Technologies*, pages 253–286. Heidelberg: Springer. quant-ph/0401019.
- Gell-Mann, M. and Hartle, J. (1990). Quantum mechanics in the light of quantum cosmology. In Zurek, W., editor, *Complexity, Entropy and the Physics of Information*, pages 425–458. Reading, MA: Addison-Wesley.
- Gell-Mann, M. and Hartle, J. (1993). Classical equations for quantum systems. *Physical Review D*, 47:3345–3382.
- Griffiths, R. B. (1984). Consistent histories and the interpretation of quantum mechanics. *Journal of Statistical Physics*, 36(1/2):219–272.
- Griffiths, R. B. (2002). Nature and location of quantum information. *Physical Review A*, 66:012311.
- Griffiths, R. B. (2003). *Consistent Quantum Theory*. Cambridge: Cambridge University Press.
- Griffiths, R. B. (2005). Channel kets, entangled states, and the location of quantum information. *Physical Review A*, 71:042337.
- Gross, D. and Eisert, J. (2006). Computational potency of quantum many-body systems. quant-ph/0609149.

- Landau, L. J. and Streater, R. F. (1993). On Birkhoff's theorem for doubly stochastic completely positive maps of matrix algebras. *Linear Algebra and its Applications*, 193:107–127.
- Minkowski, H. (1908). Raum und Zeit [Space and time]. In Einstein, A., Lorentz, H., Weyl, H., and Minkowski, H., editors, *The Principle of Relativity*, pages 72–91. London: Methuen 1923. Address delivered at at the 80th Assembly of German Natural Scientists and Physicians, Cologne, 21 September 1908.
- Müller, T. (2002). Branching space-time, modal logic and the counterfactual conditional. In Placek, T. and Butterfield, J., editors, *Non-locality and Modality*, pages 273–291. Dordrecht: Kluwer.
- Müller, T. (2005). Probability and causation. A branching space-times analysis. *British Journal for the Philosophy of Science*, 56:487–520.
- Müller, T. (2007). Branch dependence in the “consistent histories” approach to quantum mechanics. *Foundations of Physics*, 37(2):253–276. [quant-ph/0506051](https://arxiv.org/abs/quant-ph/0506051).
- Müller, T. and Placek, T. (2001). Against a minimalist reading of Bell's theorem: Lessons from Fine. *Synthese*, 128:343–379.
- Nielsen, M. A. and Chuang, I. L. (2000). *Quantum Computation and Quantum Information*. Cambridge: Cambridge University Press.
- Peres, A. (1985). Reversible logic and quantum computers. *Physical Review A*, 32:3266–3276.
- Placek, T. (2000). *Is Nature Deterministic?* Kraków: Jagiellonian University Press.
- Placek, T. (2002). Partial indeterminism is enough: a branching analysis of Bell-type inequalities. In Placek, T. and Butterfield, J., editors, *Non-locality and Modality*, pages 317–342. Dordrecht: Kluwer.
- Preskill, J. (1998). Reliable quantum computers. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 454(1969):385–410.
- Raussendorf, R. and Briegel, H. J. (2001). A one-way quantum computer. *Physical Review Letters*, 86:5188–5191.
- Shannon, C. and Weaver, W. (1949). *A Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.

- Shor, P. W. (1995). Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52:R2493–R2496.
- Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal of Computation*, 26(5):1484–1509.
- Thomason, R. H. (1970). Indeterministic time and truth-value gaps. *Theoria (Lund)*, 36:264–281.
- Weiner, M. and Belnap, N. (2006). How causal probabilities might fit into our objectively indeterministic world. *Synthese*, 149:1–36.
- Welsh, D. J. A. (1988). *Codes and Cryptography*. Oxford: Oxford University Press.