

The Search for Structure in Quantum Computation

Prakash Panangaden

School of Computer Science; McGill University and
Computing Laboratory; University of Oxford

Abstract. I give a non-comprehensive survey of the categorical quantum mechanics program and how it guides the search for structure in quantum computation. I discuss the example of measurement-based computing which is one of the successes of such an enterprise and briefly mention topological quantum computing which is an inviting target for future research in this area.

1 Introduction

Quantum computation has attracted (and repelled!) many members of the computer science community. On the one hand, people have been excited by new possibilities: cryptography based on physics rather than on unproven complexity assumptions [1], new algorithmic paradigms [2], error correction [3], solutions to hitherto “impossible” distributed computation tasks [4] and dramatic new possibilities like teleportation [5]. On the other hand, people have been disturbed by the strangeness of quantum mechanics which has rendered many of the traditional tools of theoretical computer science inapplicable.


In this paper I will attempt to convey something of the strangeness of quantum mechanics as well as some of the attempts being made to come to grips with quantum computation. The subjects of quantum algorithms and quantum information theory are flourishing and there are dramatic new results arriving at a regular pace. For the logic and semantics community it has been a rougher ride. Defining programming languages has not been routine [6–10] and there are many things that we do not understand yet. Entirely new challenges have been posed for type systems. It is only this year that we have a decent definition of weak bisimulation for a quantum process algebra. New models of computation – like measurement-based computing [11] and topological quantum computing – have emerged from the physics community which have posed challenges to the theoretical computer science community to formalize properly.

I will survey some of these developments and then pose some challenges for the future.

2 Strange features of quantum mechanics

By now most people in theoretical computer science understand the “standard” features of quantum mechanics: the state space is a Hilbert space, the evolution of a system is described by a unitary map, observables are hermitian operators (hence their eigenvalues are real) and the outcome of a measurement is probabilistic and yields one of the eigenvalues of the observable being measured. Even the originally surprising aspects of quantum mechanics are no longer surprises: systems can be in superpositions of states and measurement outcomes are not determined.

The concept of non-locality of information continues to confound many people. Even Einstein referred to this as “spooky action at a distance.” The point is that it is possible for the values of an observable to be not even defined in some states. The following thought experiment, due to Mermin [12], illustrates this dramatically. Consider the apparatus schematically shown below:



Set-up for Mermin's thought experiment

In the centre there is a source of particles that emits them in pairs travelling in opposite directions. The particles are detected by detectors that are separated far enough that signals cannot pass between them. There are two detectors each with 3 settings and 2 indicators: bulbs that flash red and green respectively. The detectors are set independently and uniformly at random. The detectors are not connected to each other or to the source.

Whatever the setting on a detector, the red or the green lights flash with equal probability, but never both at the same time. When the settings are the same the two detectors **always** agree. When the settings are different the detectors agree $\frac{1}{4}$ of the time! Why is this strange?

How could the detectors **always** agree when the settings are the same, even though the actual colour seems to be chosen at random? There must be some “hidden” property of the particles that determines which colour is chosen for each setting; the two correlated particles must be identical with respect to this property, *whether or not the switches are set the same way*. Let us write GGR mean that for the three settings, 1, 2, 3, the detectors flash green, green and red respectively for a type GGR particle. We are assuming it is meaningful to attribute properties like GGR to a particle.

Suppose that the settings are different and we have an RRG particle: then for two of the possible settings (1, 2 and 2, 1) the same colour flashes and for the other four settings the colours are different. Thus $\frac{1}{3}$ of the time the colours must match. This applies for any of the combinations: $RRG, RGR, GRR, GGR, GRG, RGG$. For particles of type RRR and GGG the colours **always** match whatever the settings. The inescapable conclusion is that *whatever the distribution of particle types* the probability that the lights match when the settings are different is at least $\frac{1}{3}$! This just ain't what we see in nature.

We made some basic assumptions about detectors:

Locality: what happens at one detector cannot alter what happens at the other,

Causality: a detector cannot predict the future sequence of particles and alter its behaviour.

No ordinary probabilistic automaton or MDP or whatever your favourite state-based model is, can reproduce the observed behaviour without breaking locality or causality. Capturing locality in an automaton means that the states of the system are the cross product of the states of each detector and the behaviour of each detector depends only on the local state.

The inequality,

$$Prob(\text{lights agree} | \text{settings different}) \geq \frac{1}{3},$$

is a simple special case of Bell's inequality. Quantum mechanics predicts that this inequality is violated. Bell's inequality has been **experimentally tested** and it is plainly violated but the experiments agree with the predictions of quantum mechanics which also predicts that the inequality is violated.

The point of this discussion is that the probabilistic nature of quantum mechanics does not arise as an abstraction of things that could be known. State is not enough to predict the outcomes of measurements; **state is enough to predict evolution to new states.**

These non-local effects are what give quantum computation its power. Teleportation is just a dramatic example of this.

3 Categorical quantum mechanics and graphical calculi

What formal techniques can be brought to bear on these kinds of systems? A key contribution of the logic and semantics community is compositionality. The whole point of denotational semantics was to provide a compositional understanding of systems. In the case of quantum mechanics we need to understand how to describe composite systems. It was known since the days of von Neumann [13] back in 1932 that the right way to combine the Hilbert spaces of two

systems is by means of the tensor product. The tensor product of Hilbert spaces is quite an elaborate construction. It requires not just the construction of the tensor product of vector spaces, but the definition of an inner product followed by a completion process which is topological in nature. Ultimately, von Neumann was unhappy with the elaborate mathematical machinery of Hilbert spaces and sought to retreat to some new fundamental logical principles for axiomatizing quantum mechanics. This led to the quantum logic programme [14] where the algebra of projection operators on a Hilbert space became the inspiration for the logic.

A huge amount of work was spent on the quantum logic enterprise [15], but in the end it is fair to say that it floundered on its inability to give a clean account of compositionality. Nevertheless logical ideas are indeed fundamental and the breakthrough idea was simple: axiomatize tensor product in the simplest way possible. This is due to Abramsky and Coecke [16] in a paper which appeared in the IEEE Symposium on Logic in Computer Science in 2004. As so often happens, categorists had invented the right notions already: monoidal categories. Though it may seem to many to not be an improvement to use fancy category theory instead of fancy functional analysis, the fact is that a purely elementary account can be given based on very simple process intuitions. A very accessible account of this viewpoint is contained in a set of lecture notes by Bob Coecke appropriately entitled, “Kindergarten Quantum Mechanics” [17].

At its most basic level then quantum mechanics is about how to hook up processes, either by connecting them in sequence or placing them side by side in parallel or combinations thereof. One can model processes as arrows in a category; the objects of the categories represent the input and output types of the processes. Categorical composition models sequential composition and the tensor product models parallel composition. Indeed one can intuit the correct axioms for tensor just from this modelling.

What is so special about quantum mechanics? Surely these are the same axioms one would use for any kind of process. One can easily whip up a model of, for example, asynchronous dataflow, as a monoidal category and indeed this has been done. Clearly monoidal categories are far too general; one needs to identify other crucial ingredients of quantum mechanics and incorporate them into the axioms. The Abramsky-Coecke paper identified duality of input and output as a crucial feature and the resulting class of categories are called by them strongly compact-closed categories. It does not matter what the algebraic axioms are because the essence of this structure is captured beautifully by a graphical calculus [18].

Graphical notions are now common in physics having been famously introduced by Feynman for keeping track of certain integrals that arise in quantum electrodynamics [19, 20]. Penrose introduced a beautiful graphical notation for tensor analysis [21] which was placed on a firm mathematical footing by Joyal and

Street [22]. I highly recommend the excellent survey by Selinger [18] in addition to the lecture notes of Coecke mentioned above.

The fundamental theorem [22] of the graphical language states that

Theorem 1. *An equation holds between terms in the morphism language of monoidal categories if and only if it holds up to planar isotopy in the graphical language.*

This means that diagrammatic manipulations can replace algebraic gymnastics. Furthermore, many equations turn out to be trivial in the graphical language.

There are two fundamental structural aspects of quantum mechanics that are captured by the categorical formalism. The first states that “objects have duals”; in categorical jargon these are called *autonomous categories*. In diagrammatic terms it means that every object A has a dual A^* and one can reverse arrows using duality. In a closed category one *bend arrows around using the duals*. This gives quantum mechanics its reversible aspect. Finally, there is a structure called a “dagger”; this is also a way of changing the direction of an arrow and it closely corresponds to the adjoint operation in linear algebra. It is the presence of this dagger structure that signals the role of complex numbers in quantum mechanics. Analogues of the fundamental theorem hold [18] for all these richer types of monoidal categories.

There are at least three important reasons for working at this level of abstractness. First, one can explore variants of quantum mechanics that are close to but not exactly the same as standard quantum mechanics. For example, the category **Rel** of sets and binary relations is an impoverished example of “toy” quantum mechanics. One can then explore what features one really needs for various phenomena to manifest themselves and thus understand what is the essence of quantum mechanics. For example, one can ask whether teleportation could be done in **Rel**; it cannot! Another striking exploration of this kind is the work by Coecke, Edwards and Spekkens [23] on formalizing a certain toy model originally due to Spekkens and showing that there is a group-theoretic reason for the difference between the two models.

Secondly, one can explore more exotic phenomena like multipartite quantum entanglement [24] or interacting quantum observables [25] from a graphical viewpoint and even find graph theoretical characterizations of these structures. As soon as one has three or more entangled states the situation becomes much more complicated. There are some preliminary hints of the role of these states in distributed computing tasks [4] but clearly much remains to be understood and structural understanding will guide the way.

Finally, one can address foundational questions of quantum information and quantum mechanics. In very striking recent work Abramsky [26] has analyzed the issue of hidden variables in the toy relational model and has shown that

many of the no-go theorems survive even when one has only a “possibilistic” view of nondeterminism.

4 Measurement-based computing

I now turn to a new computational model and analyze it from the viewpoint of theoretical computer science. Traditionally, the main framework to explore quantum computation has been the circuit model [27], based on unitary evolution. This is very useful for algorithmic development and complexity analysis [28]. There are other models such as quantum Turing machines [29] among a variety of others. They are all proved to be equivalent from the point of view of expressive power. For higher-order sequential programming we have the typed λ -calculus which stands out as *the* canonical programming language but there is no such language for quantum computation.

Recently physicists have introduced novel ideas based on the use of measurement and entanglement to perform computation [30, 11, 31]. This is very different from the circuit model where measurement is done only at the end to extract classical output. In measurement-based quantum computation the main operation to manipulate information and control computation is measurement. This is surprising because measurement creates indeterminacy, yet it is used to express deterministic computation defined by a unitary evolution.

The idea of computing based on measurements emerged from the teleportation protocol [5]. The goal of this protocol is for an agent to transmit an unknown qubit to a remote agent without actually sending the qubit. This protocol works by having the two parties share a maximally entangled state called a Bell pair. The parties perform *local* operations – measurements and unitaries – and communicate only classical bits. Remarkably, from this classical information the second party can reconstruct the unknown quantum state. In fact one can actually use this to compute via teleportation by choosing an appropriate measurement [30]. This is the key idea of measurement-based computation.

It turns out that the above method of computing is actually universal. This was first shown by Gottesman and Chuang [30] who used two-qubit measurements and given Bell pairs. The one-way computer was then invented by Raussendorf and Briegel [11, 32] which used only single-qubit measurements with a particular multi-party entangled state called the cluster state.

The computation proceeds in a sequence of phases; in the first phase a collection of qubits are set up in a standard entangled state. Then measurements are applied to individual qubits and the outcomes of the measurements may be used to determine further adaptive measurements. Finally – again depending on measurement outcomes – local unitary operators, called corrections, are applied to some qubits; this allows the elimination of the indeterminacy introduced by

measurements. The phrase “one-way” is used to emphasize that the computation is driven by irreversible measurements.

There are at least two reasons to take measurement-based models seriously: one conceptual and one pragmatic. The main pragmatic reason is that the *one-way* model is believed by physicists to lend itself to easier implementations [33–35]. Physicists have investigated various properties of the cluster state and have accrued evidence that the physical implementation is scalable and robust against decoherence [36]. Conceptually the measurement-based model highlights the role of entanglement and separates the quantum and classical aspects of computation; thus it clarifies, in particular, the interplay between classical control and the quantum evolution process.

When this model was first presented it was couched in the language of Hamiltonians and evolution of quantum states. The design of even simple gates seemed magical and an exercise in combinatorial ingenuity. Most importantly, the “proof” of universality consisted in showing that the basic gates of the circuit model could be implemented in the one-way model with the implicit understanding that any network could then be encoded. What was missing was a *compositional translation* and a proof that the *semantics* of the circuit was preserved.

Our approach to understanding the structural features of measurement-based computation was to develop a formal calculus [37]. One can think of this as an “assembly language” for measurement-based computation. It was the first programming framework specifically based on the one-way model. In our paper we developed a language for programs (we called them “patterns”) as sequences of entanglements, measurements, and local corrections. We gave a careful treatment of the composition and tensor product (parallel composition) of programs and we have denotational semantics and operational semantics for these programs. In this framework we were able to give a proof of universality. In fact, we were able to modify the framework in apparently small ways but these had the effect of greatly simplifying the implementations of circuits. More precisely we had an extended notion of pattern, where inputs and outputs may overlap in any way one wants them to, and this results in more efficient – in the sense of using fewer qubits – implementations of unitaries. Specifically, our universal set consists of patterns using only 2 qubits. From it we obtained a 3 qubit realization of the R_z rotations and a 14 qubit realization for the controlled- U family: a significant reduction over the hitherto known implementations [38].

However, there were more benefits to be gained from the exploration of this structural view of measurement-based computing. We introduced a *calculus* of patterns based on the special algebraic properties of the entanglement, measurement and correction operators. These allowed local rewriting of patterns and we showed that this calculus is sound in that it preserves the interpretation of patterns. Most importantly, we derived from it a simple algorithm by which any general pattern can be put into a standard form where entanglement is done first, then measurements, then corrections. We call this *standardization*.

The consequences of the existence of such a procedure are far-reaching. Since entangling comes first, one can prepare the entire entangled state needed during the computation right at the start: one never has to do “on the fly” entanglements. Furthermore, the rewriting of a pattern to standard form reveals parallelism in the pattern computation. In a general pattern, one is forced to compute sequentially and to strictly obey the command sequence, whereas, after standardization, the dependency structure is relaxed, resulting in lower computational depth complexity [39].

Perhaps the most striking development in the theory of measurement-based computing is the discovery of the concept of *flow* by Danos and Kashefi [40]. A variety of methods for constructing measurement patterns had already been proposed that guarantee determinism by construction. They introduced a graph-theoretic condition on the states used in measurement patterns that guarantees a strong form of deterministic behavior for a class of one-way measurement patterns defined over them. Remarkably, their condition bears only on the geometric structure of the entangled graph states. This condition singles out a class of patterns with flow, which is stable under sequential and parallel compositions and is large enough to realize all unitary and unitary embedding maps.

Patterns with flow have interesting additional properties. First, they are uniformly deterministic, in the sense that no matter what the measurements are made, there is a set of corrections, which depends only on the underlying geometry, that will make the global behaviour deterministic. Second, all computation branches have equal probabilities, which means in particular, that these probabilities are independent of the inputs, and as a consequence, one can show that all such patterns implement unitary embeddings. Third, a more restricted class of patterns having both flow and reverse flow supports an operation of adjunction, corresponding to time-reversal of unitary operations. This smaller class implements all and only unitary transformations.

In the categorical quantum framework Coecke and Duncan [25] have looked at interacting quantum observables in a diagrammatic formalism. There is a very pleasing encoding of the one-way model into this framework and many of the algebraic equations of the one-way model can be done by graphical manipulations. It would be fascinating to understand flow and its relation to causality in this way.

5 Topological quantum computing

I would like to close with a brief mention of a wide open area: topological quantum computing. In quantum computation one is required to make excruciatingly precise manipulations of qubits while preserving entanglement when all the while the environment is trying to destroy entanglement. A very novel suggestion by Kitaev [41] proposes the use of a new type of hypothetical particle called an anyon [42, 43] which has a topological character.

The mathematics and physics of anyons probe the most fundamental principles of quantum mechanics. They involve a fascinating mix of experimental phenomena (the fractional quantum Hall effect), topology (braids), algebra (Temperley-Lieb algebra, braid group and category theory) and quantum field theory. Because of their topological nature, it is hoped that one can use them as *stable* realizations of qubits for quantum computation, as proposed originally by Kitaev. The idea of topological quantum computation has been actively pursued by Freedman et al. [44, 45].

There is rich algebraic structure to be understood and, as with the measurement-based model, we need a computational handle on this. We have nothing like a calculus or a standardization theorem. What is clear is that the basic interactions of the anyons can only be expressed in the categorical language. One needs a rather rich kind of categorical structure called a modular tensor category [46]. An expository account of this subject is given in [47]. Understanding topological quantum computing from the viewpoint of computational structure remains a big open problem.

Acknowledgments

I would like to thank McGill University and EPSRC for its generous support during my sabbatical leave, NSERC (Canada) and the Office of Naval Research (USA) for funding this research and the Computing Laboratory of the University of Oxford for its hospitality. I have benefitted from discussions with Samson Abramsky, Bob Coecke, Vincent Danos, Ross Duncan, Julia Evans, Elham Kashefi, Eric Paquette and Jamie Vicary.

References

1. Bennett, C.H., Brassard, G.: Quantum cryptography: Public-key distribution and coin tossing. In: Proceedings of IEEE International Conference on Computers, Systems and Signal Processing, Bangalore, India. (December 1984) 175–179
2. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In Goldwasser, S., ed.: Proc. 35th Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press (1994) 124–134
3. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Computing* **26** (1997) 1484–1509
4. D’Hondt, E., Panangaden, P.: The computational power of the W and GHZ states. *Quantum Information and Computation* **6**(2) (2006) 173–183
5. Bennett, C.H., Brassard, G., Crepeau, C., Josza, R., Peres, A., Wootters, W.: Teleporting an unknown quantum state via dual classical and epr channels. *Phys. Rev. Lett.* **70** (1993) 1895–1899
6. Gay, S.: Quantum programming languages: Survey and bibliography. *Bulletin of the EATCS* **86** (June 2005) 176–196

7. Selinger, P.: Towards a quantum programming language. *Mathematical Structures in Computer Science* **14**(4) (2004) 527–586
8. Selinger, P.: A brief survey of quantum programming languages. In: *Proceedings of the 7th International Symposium on Functional and Logic Programming*. Number 2998 in *Lecture Notes In Computer Science*, Springer-Verlag (2004) 1–6
9. Selinger, P., Valiron, B.: Quantum lambda calculus. In Gay, S., Mackie, I., eds.: *Semantic Techniques in Quantum Computation*. Cambridge University Press (2009) to appear.
10. van Tonder, A.: A lambda calculus for quantum computation. *Siam Journal on Computing* **33**(5) (2004) 1109–1135
11. Raussendorf, R., Briegel, H.J.: A one-way quantum computer. *Phys. Rev. Lett.* **86** (May 2001) 5188–5191
12. Mermin, D.: *Boojums all the way through*. Cambridge University Press (1990)
13. von Neumann, J.: *Mathematisch Grunlagen der Quantenmechanik*. Springer-Verlag (1932) English translation, 1955, Princeton University Press.
14. Birkhoff, G., von Neumann, J.: The logic of quantum mechanics. *Annals of Mathematics* **37**(4) (1936) 823–843
15. Piron, C.: *Foundations of quantum physics*. W. A. Benjamin (1976)
16. Abramsky, S., Coecke, B.: A categorical semantics of quantum protocols. In: *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science: LICS 2004*, IEEE Computer Society (2004) 415–425
17. Coecke, B.: Kindergarten quantum mechanics. Available on the ArXiv [quant-ph/0510032](https://arxiv.org/abs/quant-ph/0510032) (2005)
18. Selinger, P.: A survey of graphical languages for monoidal categories. In: *New Structures for Physics*. Springer-Verlag (2010) 289–356
19. Feynman, R.P.: The theory of positrons. *Physical Review* **76** (1949) 749–759
20. Feynman, R.P.: The space-time approach to quantum electrodynamics. *Physical Review* **76** (1949) 769–789
21. Penrose, R.: Applications of negative dimensional tensors. In Welsh, D.J.A., ed.: *Combinatorial Mathematics and its Applications*. Academic Press (1971)
22. Joyal, A., Street, R.: The geometry of tensor calculus. *Advances in Mathematics* **88** (1991) 55–112
23. Coecke, B., Edwards, B., Spekkens, R.: The group theoretic origin of non-locality for qubits. Technical Report RR-09-04, OUCL (2009)
24. Coecke, B., Kissinger, A.: The compositional structure of multipartite quantum entanglement. In: *ICALP (2)*. (2010) 297–308
25. Coecke, B., Duncan, R.: Interacting quantum observables. In: *ICALP (2)*. (2008) 298–310
26. Abramsky, S.: Relational hidden variables and non-locality. arXiv:1007.2754 (July 2010)
27. Deutsch, D.: Quantum computational networks. *Proc. Roy. Soc. Lond A* **425** (1989)
28. Bernstein, E., Vazirani, U.: Quantum complexity theory. *SIAM Journal of Computing* **5**(26) (1997)
29. Deutsch, D.: Quantum theory, the Church-Turing Principle and the universal quantum computer. *Proc. Roy. Soc. Lond. A* **400** (1985) 97
30. Gottesman, D., Chuang, I.L.: Quantum teleportation is a universal computational primitive. *Nature* **402** (1999)
31. Nielsen, M.A.: Universal quantum computation using only projective measurement, quantum memory, and preparation of the 0 state. *Physical Review A* **308** (2003)

32. Raussendorf, R., Browne, D.E., Briegel, H.J.: Measurement-based quantum computation on cluster states. *Phys. Rev. A* **68**(2) (2003) 022312
33. Nielsen, M.A.: Optical quantum computation using cluster states. *Physical Review Letters* **93** (2004) quant-ph/0402005.
34. Childs, A.M., Leung, D.W., Nielsen, M.A.: Unified derivations of measurement-based schemes for quantum computation. *Physical Review A* **71** (2005) quant-ph/0404132.
35. Browne, D.E., Rudolph, T.: Resource-efficient linear optical quantum computation. *Physical Review Letters* **95** (2005) quant-ph/0405157.
36. Hein, M., Eisert, J., Briegel, H.J.: Multi-party entanglement in graph states. *Physical Review A* **69** (2004) quant-ph/0307130.
37. Danos, V., Kashefi, E., Panangaden, P.: The measurement calculus. *Journal Of The Association Of Computing Machinery* **52**(2) (April 2007) article 8.
38. Vincent Danos, E.K., Panangaden, P.: Parsimonious and robust realizations of unitary maps in the one-way model. *Physical Review A* **72** (Dec 2005) 064301
39. Broadbent, A., Kashefi, E.: Parallelizing quantum circuits. *Theoretical Computer Science* **410**(26) (2009) 2489–2510
40. Danos, V., Kashefi, E.: Determinism in the one-way model. *Physical Review A* **74**(5) (2006) 6 pages
41. Kitaev, A.: Fault-tolerant quantum computation by anyons. *Ann. Phys.* **303**(1) (2003) 3–20
42. Wilczek, F.: Magnetic flux, angular momentum, and statistics. *Phys. Rev. Lett.* **48**(17) (April 1982) 1144–1146
43. Wilczek, F.: Quantum mechanics of fractional-spin particles. *Phys. Rev. Lett.* **49**(14) (Oct 1982) 957–959
44. Freedman, M.H., Larsen, M., Wang, Z.: A modular functor which is universal for quantum computation. *Communications in Mathematical Physics* **227**(3) (2002) 605–622
45. Freedman, M.H., Kitaev, A., Larsen, M., Wang, Z.: Topological quantum computing. *Bulletin of the AMS* **40**(1) (2003) 31–38
46. Bakalov, B., Kirillov, A.: Lectures on tensor categories and modular functors. American Mathematical Society in University Lecture Series (2001)
47. Panangaden, P., Paquette, E.: A categorical presentation of quantum computation with anyons. In: *New Structures for Physics*. Springer-Verlag (2010) 983–1026