Categorical Quantum Computing: Diagrammatic Reasoning for Quantum Algorithms

James G. Morley*
Centre for Doctoral Training in Quantum Technologies, University College, London WC1E 6BT
(Dated: May 6, 2016)

Algorithms for quantum computing are an important cornerstone in the field of quantum technologies. However the language we use to describe them can be clunky and awkward - expositions of quantum protocols are often interspersed with explanatory sentences. This inhibits the understanding of such protocols and also the designing of new protocols. In this case study report I discuss a new approach to the language of quantum processes that uses intuitive diagrammatic reasoning backed up with the rigorous mathematics of category theory.

I. INTRODUCTION

Shor's factoring algorithm [1] kick-started the field of quantum computing algorithms [2] and sparked a worldwide program of research in quantum technologies that is still growing more than 20 years on - the European commission announced last month plans for a € 1 billion investment in a 'large-scale EU-wide quantum technologies flagship' [3, 4].

Quantum algorithms are still important today; they promise huge speedups over classical algorithms for important computational tasks [5]. However, just like with classical algorithms, it is difficult to design new quantum algorithms. It is also difficult to understand how quantum algorithms work - there has been much discussion over precisely which feature of quantum theory gives quantum algorithms their edge [6].

This case study is about designing new formalisms for quantum theory that will give us a clearer and more natural language for thinking about quantum computing algorithms, as well as other topics in quantum theory more generally.

Part of the difficulty in analyzing algorithms lies in the established formalism of quantum mechanics. The three main paradigms for studying quantum computing are the circuit or gate-based model [7], quantum walks [8] and adiabatic quantum computation [9]. The standard presentations of each look very different, yet the same quantum theory underlies them all - they all consist of quantum and classical systems undergoing quantum and classical maps. The same can be said for the quantum teleportation protocol [10], and quantum communication protocols generally.

Established quantum theory is grounded in von Neumann's Hilbert space formalism [11, 12]. One of the first critics of von Neumann's formalism was, surprisingly, von Neumann himself! In a letter to Garrett Birkhoff in 1935, von Neumann said

> *"I would like to make a confession that may seem immoral: I do not believe absolutely in Hilbert space no more* (sic)*."*

What brought von Neumann to this belief (or lack of) was the abstract nature of the mathematical definitions that lay at the heart of his formalism. True to his word, he devoted much of his later career to an, ultimately unsucessful, pursuit of a better formulation [13].

To be more precise, two features of Hilbert space in particular make quantum algorithm design difficult. Firstly, the foundations of von Neumann's quantum mechanics are couched in abstract mathematical definitions - there is no *operational* meaning to them. This contrasts with the other great physical theory, relativity, whose ideas are built from the very operational concepts of moving rods and clocks and using them to measure time and length. Secondly, the theory is very 'low-level': many of the symbols we write in the algebra of Hilbert spaces are required purely for mechanistic book-keeping, they don't tell us anything about what's really happening. A strong analogy can be made here with programming languages: computer programs written in low-level programming languages are much longer than those written in high-level languages for which the mechanistic book-keeping is implicit. For quantum computing algorithms a higher level of language would be invaluable [14].

This is where the word 'categorical' comes in. Over the past three decades or so work has been underway in defining new quantum theories that have *operational* foundations and can be described in a *higher-level* language. The mathematics underlying this work is called **category theory**. Category theory is the mathematical theory of systems ('objects' in category theory parlance) and the processes ('morphisms') that they undergo [15]. This extremely general mathematics has several nice features for quantum theory: it is process-oriented and so inherently operational, it deals very naturally with composite systems (entanglement is fundamental to quantum theory [16] and so composite systems should be fundamental in our formalism), and it allows us to build up quantum theories axiomatically - something that comes in very useful in foundational questions in quantum theory.

Perhaps the most appealing aspect of category theory is that for many of the quantum theories built from it a *picture calculus* is permitted. That is, calculations

can be represented using diagrams, and these diagrams can be manipulated to prove new equalities in a rigorous fashion. Moreover these picture calculuses can be 'complete': meaning that anything provable in the theory can be proved with graphical manipulations alone. This is astonishing: once the rules for creating and manipulating diagrams are established, the algebra can be thrown away!

In the next section I will give an overview of how category theory is used to build high-level, operational quantum theories with composition of systems built in naturally. In sections 3-5 I will present recent progress in developing such formalisms, and using them to better understand quantum algorithms and quantum theory more widely. Section 6 concludes the report with speculations on future prospects of diagrammatic reasoning for quantum algorithms.

II. BUILDING DIAGRAMMATIC THEORIES

The use of diagrams to aid calculation isn't a new concept, it emerges naturally in all areas of science and mathematics. Perhaps the most famous example in physics is the Feynman diagram. Attempts to imbue these diagrams with a rigorous mathematical meaning really began with a study of Feynman diagrams and a collection of other diagrammatic calculational aids in 1991 [17]. This is when the power of category theory in formulating diagrammatic theories first came to light.

Soon after this result, work began in earnest, and has continued ever since, to try and formulate a quantum mechanics imbued with the rigour of category theory [18–22].

So how do we go about building a diagrammatic theory using category theory? Category theory is the study of processes, and so category theories are sometimes referred to as 'process theories'. The most generic process theory within this framework is the theory of single 'objects' $A, B, C \ldots$ (that is, systems: A might be a single qubit, B a qubit plus a classical bit, etc.) undergoing 'morphisms' $f, g, h \ldots$ (that is, processes: f might be the morphism that takes a qubit and appends a classical bit, thereby mapping object A to object B). We can represent processes in this theory with diagrams in which objects are represented by solid lines or 'wires', and morphisms as boxes with input and output wires. Figure 1(a) shows such a diagram: conventionally time progresses up the page, so this diagram shows the morphism f that maps object A to object B. As well as objects and morphisms we require a rule for composing morphisms sequentially, usually denoted '$\circ$'. Figure 1(b) shows a diagram for the composite morphism $g \circ f$, where g is a morphism that maps object B to object C.

From this general starting point we add new elements of our diagrammatic theory to enable it to describe more

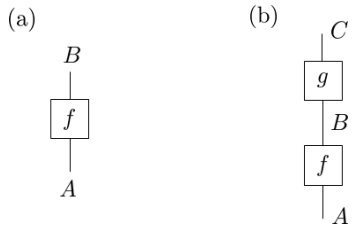

FIG. 1. Elementary category-theoretic diagrams. As drawn here time progresses upwards. (a) morphism f maps object A to B, that is $f : A \mapsto B$. (b) $g : B \mapsto C$, so the composite $g \circ f : A \mapsto C$.

complex processes. In this way we can build up theories axiomatically. For a quantum theory we need, at the minimum, a 'parallel composition' or tensor product $\otimes$ that allows us to consider composite objects and morphisms, and also a swap operation for exchanging the position of two objects. For a more thorough description, see Bob Coecke's paper "Introducing Categories to the Practicing Physicist" [23].

Categorical diagrammatic theories give us a high-level language for describing complex sequences of processes between quantum systems. While this can't describe anything not already present in the Hilbert space formalism of quantum mechanics, it *can* expose the structure of quantum algorithms and other complicated quantum protocols. Category theory provides the background mathematical framework for building these process theories: from this new diagrammatic theories can be crafted, and from these complicated quantum computing and communication protocols can be analyzed with reference only to the high-level diagrammatic language provided.

This isn't to say that categorical diagrammatic theories should supplant Hilbert space quantum mechanics entirely - sometimes it will be appropriate to deal with the low-level explicit detail (this is exactly what happens in programming languages). However, when used according to their strengths, the diagrammatic theories discussed here can and do facilitate new insights.

III. TOPOLOGY OF QUANTUM ALGORITHMS

In 2013 a result was published that exemplifies the use of high-level categorical quantum theories to gain new insight into quantum computing algorithms [24]. Using such a theory, the canonical Deutsch-Jozsa [25], single-shot Grover [26] and Hidden-subgroup [27] algorithms were formulated diagrammatically and analyzed. The resulting process diagrams are shown in figure 2.

Recalling that time progresses upwards in these diagrams, we see that each algorithm is separated into preparation (1), dynamics (2) and measurement (3) phases. The dynamics phase consists of a unitary oracle operations, common to each algorithm. The oracle

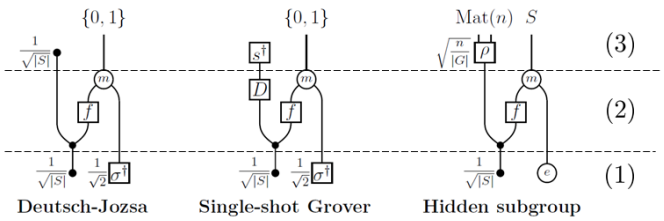

FIG. 2. Categorical diagrammatic representations of the Deutsch-Jozsa, single-shot Grover and hidden subgroup algorithms [24]. Each algorithm consists of three steps, indicated by dashed lines (these form no part of the diagrams themselves). These are the *preparation* phase (1), the *dynamics* phase (2) and the *measurement* phase (3).

maps a state of the form $|i\rangle\,|j\rangle$ to $|i\rangle\,|f(i)+j\rangle$, where '+' is the relevant addition operation and the function f characterized in the task the algorithm is attempting to perform. The diagrammatic form of unitary oracles is key to each representation in figure 2, and is explored in more detail in [28].

These diagrammatic representations expose these well-studied algorithms in a new light: the diagrams clearly have a common topological structure. Using this insight new proofs of correctness were found for the three algorithms, and generalizations of each algorithm were presented. The generalization of the single-shot Grover algorithm was completely new. Generalizations for the Deutsch-Jozsa and Hidden subgroup algorithms already existed [27, 29, 30], however the high-level description permitted by this formulation allowed for an improved analysis.

IV. ZX-CALCULUS

'ZX-Calculus' is the name given to a diagrammatic process theory for pure state qubit quantum mechanics with post-selected measurements. It was introduced in [18, 31] and refined in [32]. By virtue of its category-theoretic grounding it is more flexible than the standard circuit notation used in gate-based quantum computing [7], and it is more fundamental too - the elementary physical gates of quantum computing are seen in ZX-calculus to be built from more fundamental, though individually unphysical, components [33].

While developing the ZX-calculus has taken some time, the formalism itself is straightforward to learn. Its diagrams are built from a set of elementary building blocks that are straightforwardly defined. Particularly useful for learning ZX-calculus is a built-in 'usual interpretation of' mapping from ZX-calculus diagrams to their 'usual interpretation' in terms of the familiar Dirac notation (up to normalization factors). For example, the usual interpretation of a wire, denoted $[\![\ |\ ]\!]$, is the identity operation:

$$[\![\ |\ ]\!] = |0\rangle\langle 0| + |1\rangle\langle 1|\,. \tag{1}$$

The most important building blocks of ZX-calculus are the eponymous Z and X 'phase spiders' (diagrams taken from [32]):

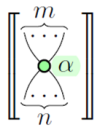

$$= |0\rangle^{\otimes m}\langle 0|^{\otimes n} + e^{i\alpha}\,|1\rangle^{\otimes m}\langle 1|^{\otimes n}\,, \tag{2}$$

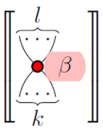

$$= |+\rangle^{\otimes l}\langle +|^{\otimes k} + e^{i\beta}\,|-\rangle^{\otimes l}\langle -|^{\otimes k}\,, \tag{3}$$

where $|+\rangle$ and $|-\rangle$ are the eigenstates of the single-qubit Pauli-X operator. These can take arbitrary numbers of input and output qubits, and act as generalized phase gates in complementary bases usually referred to as Z and X after the standard Pauli-Z and Pauli-X bases of a qubit [7].

Adding diagram elements for transformations between the bases, a SWAP operation, creation and measurement of a Bell state, and for the number 1/2 we have a full list of the building blocks for the ZX-calculus.

With this fairly simple list much can be achieved. The biggest early success of ZX-calculus was its use in formulating a high-level description of measurement based quantum computing (MBQC) [34]. In MBQC a complex entangled many-qubit state is prepared, after which single qubit measurements are performed. This, surprisingly, results in a computationally useful state. When formulated in the language of ZX-calculus, a little diagrammatic reasoning makes clear the relationship between the single qubit measurements and the effects they have on the final system state.

The full scope of ZX-calculus is an open question. It has been shown to be complete for stabilizer quantum mechanics on pure qubit states [35, 36], and for an approximately universal gate on a single qubit [32]. Proving approximate universality for larger systems is a topic of current research. Whatever the outcome of those investigations, the ZX-calculus should certainly be taken seriously as an alternative to conventional quantum circuit diagrams.

V. QUANTUM FOUNDATIONS - SPEKKEN'S TOY MODEL

Exactly what it is that gives quantum algorithms their speedup over their classical counterparts is an important question [6]. Very closely related is the more fundamental question of 'what is it that makes quantum mechanics

quantum?' - is absolutely all of the machinery required, or does it boil down to a small number of key ingredients? In 2004 Robert Spekkens presented a toy theory that went some way to answering that question [37]. With a very simple theory, thereafter known as 'Spekken's toy model', many of the phenomena normally associated exclusively with quantum mechanics could be reproduced, including for example incompatible observables, teleportation, and no-cloning. One important difference however between the toy model and quantum theory is that the toy model is a *local* theory - the Bell inequalities aren't violated [38].

In 2011 Coecke, Edwards and Spekkens took a categorical diagrammatic approach to try to pin down the difference between quantum theory and the toy model even further [39]. They defined a diagrammatic theory for the toy model, which they dubbed **Spek**, and for stabilizer quantum mechanics (a restricted version of full quantum theory [40]) which they called **Stab**. This provided them with an equivalent diagrammatic machinery with which to analyze the two theories.

It was found that while the *physical* reason for the difference in the two theories was that **Spek** is a local theory whereas **Stab** is not, it can further be said that the *structural* difference lay in the 'phase group' of individual systems. Found in either theory are three pairs of orthogonal state for an individual system that can be arranged along three axes - the Bloch sphere arrangement for **Stab** and an analogous arrangement for **Spek**. The 'phase group' of the theories is the group structure of the equatorial states, as defined by the transformations allowed in the theory. The phase group of **Spek** is $\mathbb{Z}_4$ and the phase group of **Stab** is $\mathbb{Z}_2 \times \mathbb{Z}_2$, where $\mathbb{Z}_n$ is the group of $2\pi/n$ planar rotations. This result pins down more precisely the role non-locality plays in quantum theory.

The graphical calculus for Spekken's toy model has since been developed further and shown to be complete [36] - that is, any equality derivable in the theory can be derived with diagrammatic reasoning alone. This opens the door to further investigations into the nature of non-locality in quantum theory, with potential implications for quantum protocol design that takes advantage of non-local effects, including computing and communication.

VI. CONCLUSIONS

To summarize, the body of work discussed here shows us that we can build graphical theories of quantum processes that are mathematically rigorous, founded on operational ideas and that allow high-level descriptions of complicated processes. These complicated processes can be represented with two-dimensional diagrams rather than the normal one-dimensional representation of lines of algebra. The evolution of composite systems is inherently a process in at least two dimensions - one dimension for time, and at least one for the arrangement of the subsystems. This is essentially why a diagrammatic description is high-level: the mechanistic book-keeping normally required in this compression to one dimension becomes superfluous and can easily, and rigorously, be made implicit.

Future Work

The usefulness of the diagrammatic theories discussed here is only just beginning to be tapped. Substantial work has already been completed in developing suitable graphical theories for quantum computing algorithms using the tools of category theory, and these theories are now ready to be used themselves as useful tools. There is of course further work to be done in theory development; in the near future we can hope to see results regarding the question of whether the ZX-calculus is complete for an approximately universal gate set on multiple qubits. General quantum theory can also be formulated as a graphical theory in this way, the details of this are to be published later this year in an upcoming textbook on diagrammatic reasoning and quantum theory [41].

Diagrammatic theories provide high-level descriptions of quantum processes that are particularly useful for complicated protocols involving multiple interacting quantum systems. Uptake of the graphical approach to analyzing quantum computing algorithms has so far been limited, and so there are a limited number of results published so far. The causation here goes both ways however - the lack of many results, or of one very significant result, leaves many unconvinced that switching over to a new framework for quantum algorithms is worth the effort. A significant result would change this, so the future of the approach described here hinges on whether diagrammatic reasoning can directly aid in pushing back the boundaries of knowledge in a significant jump that reveals something truly *new*.

ACKNOWLEDGEMENTS

I would like to thank Dan Browne, Bob Coecke and Stefano Gogioso for helpful discussions while carrying out this investigation.

* james.morley.15@ucl.ac.uk

[1] P. W. Shor. *SIAM J. Comput.*, 26:1484, 1997 (preliminary version in FOCS 1994).
[2] A. M. Childs and W. van Dam. *Rev. Mod. Phys.*, 82:1, 2010.

[3] Digital single market digitising european industry questions & answers. http://europa.eu/rapid/press-release_MEMO-16-1409_en.htm. Accessed 29/04/2016.

[4] Europe plans giant billion-euro quantum technologies project. doi:10.1038/nature.2016.19796. Accessed 29/04/2016.

[5] A. Montenaro. *npj Quantum Information*, 2:15023, 2016.

[6] A. Galindo and M. A. Martín-Delgado. *Rev. Mod. Phys.*, 74:347, 2002.

[7] M. A. Nielsen and I. L. Chuang. *Quantum Computing and Quantum Information*. CUP, Cambridge, 2000.

[8] A. M. Childs. *Phys. Rev. Lett.*, 102:180501, 2009.

[9] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser. *arXiv:quant-ph/0001106*, 2000.

[10] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, Peres A., and W. K. Wootters. *Phys. Rev. Lett.*, 70:1895, 1993.

[11] J. von Neumann. *Math. Ann.*, 102:49–131, 1927.

[12] L. van Hove. *Bull. Amer. Math. Soc.*, 64:95–99, 1958.

[13] G. Birkhoff and J. von Neumann. *Annals of Mathematics*, 37:823–843, 1936.

[14] B. Coecke. *arXiv:quant-ph/0510032*, 2005.

[15] Saunders Mac Lane. *Categories for the working mathematician*. Springer, New York, 2nd edition, 1998.

[16] E. Schrödinger. Discussion of probability relations between separated systems. In *Proceedings of the Cambridge Philosophical Society 31*, pages 555–563. 1935.

[17] A. Joyal. *Advances in Math.*, 88:55–112, 1991.

[18] B. Coecke and R. Duncan. *New J. Phys.*, 13:043016, 2011.

[19] B. Coecke, E. O. Paquette, and D. Pavlovic. In *Semantic Techniques in Quantum Computation*, pages 29–69. Cambridge University Press, 2010.

[20] B. Coecke and S. Perdrix. In *Proceedings of the 19th EACSL Annual Conference on Computer Logic*. Number 6247 in Lecture Notes in Computer Science, 2010.

[21] D. J. Moore. *Helv. Phys. Acta*, 68:658–678, 1995.

[22] C. J. Wood, J. D. Biamonte, and D. G. Cory. *Quant. Inf. Comp.*, 15:0579–0811, 2015.

[23] B. Coecke. Introducing categories to the practicing physicist. *What is category theory? Advanced Studies in Mathematics and Logic*, 30:45–74, 2008. URL http://arxiv.org/abs/0808.1032.

[24] J. Vicary. In *Proceedings of the 28th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 93–102. 2013.

[25] D. Deutsch and R. Jozsa. *Proc. R. Soc. London, Ser. A*, 439:553, 1992.

[26] L. Grover. *Phys. Rev. Lett.*, 79:325–328, 1997.

[27] S. Hallgren, A. Russell, and A. Ta-shma. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, pages 916–934. 2000.

[28] W. Zeng and J. Vicary. *EPTCS*, 172:270–284, 2014.

[29] P. Hø yer. *Phys. Rev. A*, 59:3280–3289, 1999.

[30] A. J. Duncan, M. Batty, and S. L. Braunstein. *J. Comp. and Math.*, 9:40–63, 2006.

[31] B. Coecke and R. Duncan. Automata, languages and programming. volume 5126, pages 298–310. Springer Berlin Heidelberg, Berlin, Heilberg, 2008.

[32] M. K. Backens. *Completeness and the ZX-Calculus*. PhD thesis, University of Oxford, 2015.

[33] R. Duncan and S. Perdrix. In *Lecture Notes in Computer Science*, pages 167–177. 2009.

[34] R. Duncan and S. Perdrix. Automata, languages and programming. volume 6199, pages 285–296. Springer Berlin Heidelberg, Berlin, Heilberg, 2010.

[35] D. Gottesman and I. L. Chuang. *Nature*, 402:390–393, 1999.

[36] M. Backens. *New J. Phys.*, 16:093021, 2014.

[37] R. W. Spekkens. *Phys. Rev. A*, 75:032110, 2004.

[38] J. S. Bell. *Physics*, 1:195–200, 1964.

[39] B. Coecke, B. Edwards, and R. W. Spekkens. *Electrical Notes in Th. Comp. Sci.*, 270:15–36, 2011.

[40] D. Gottesman. *Stabilizer Codes and Quantum Error Correction*. PhD thesis, California Institute of Technology, 1997.

[41] B. Coecke and A. Kissinger. *Picturing Quantum Processes*. Cambridge University Press, (to appear).