

Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy

BY MICHAEL J. BREMNER^{1,*}, RICHARD JOZSA² AND DAN J. SHEPHERD³

¹*Institut für Theoretische Physik, Leibniz Universität Hannover,*

Appelstrasse 2, Hannover 30167, Germany

²*DAMTP, Centre for Mathematical Sciences, University of Cambridge,*

Wilberforce Road, Cambridge CB3 0WA, UK

³*CESG, Hubble Road, Cheltenham GL51 0EX, UK*

We consider quantum computations comprising only commuting gates, known as IQP computations, and provide compelling evidence that the task of sampling their output probability distributions is unlikely to be achievable by any efficient classical means. More specifically, we introduce the class post-IQP of languages decided with bounded error by uniform families of IQP circuits with post-selection, and prove first that post-IQP equals the classical class PP. Using this result we show that if the output distributions of uniform IQP circuit families could be classically efficiently sampled, either exactly in total variation distance or even approximately up to 41 per cent multiplicative error in the probabilities, then the infinite tower of classical complexity classes known as the polynomial hierarchy would collapse to its third level. We mention some further results on the classical simulation properties of IQP circuit families, in particular showing that if the output distribution results from measurements on only $O(\log n)$ lines then it may, in fact, be classically efficiently sampled.

Keywords: quantum computational complexity; quantum simulation; polynomial hierarchy

1. Introduction

From a pragmatic point of view, the field of quantum computing is driven by the expectation that quantum algorithms can offer some computational complexity benefits transcending the possibilities of classical computing. But this expectation can be challenged both theoretically and experimentally: (i) there is yet no theoretical proof that any quantum algorithm outperforms the best classical algorithm for the task, in the standard computational setting of polynomial versus exponential running time (without inclusion of further constructs, such as use of oracles, or consideration of distributed computing and the role of communication; in both these scenarios there are indeed proofs of exponential complexity benefits); (ii) experimentally there are well-documented difficulties associated with building a quantum computer that is suitably fault tolerant and sufficiently scalable to manifestly demonstrate a complexity benefit.

*Author for correspondence (michael.bremner@itp.uni-hannover.de).

However, both (i) and (ii) can, to some extent, be redressed by further examination: the criticism in (i) can be attributed to limitations of *classical* complexity theory—we do have interesting quantum algorithms (such as Shor’s factoring algorithm) for problems widely believed to be classically hard but there is no proof of the latter. Proof of classical hardness is a notoriously difficult issue (see the famous P versus NP question) and it has become popular to resort to providing only evidence of hardness, as follows: we prove that if a certain problem were classically easy then this would entail consequences that are highly implausible (although also generally unproven), e.g. collapse of an entire complexity class (such as entailing that $P = NP$). For (ii) we could seek to devise a computational task that, on the one hand, is expected to be classically hard (as above) yet, on the other hand, can be implemented using suitably simple (subuniversal) quantum computational elements that are especially easily or fault-tolerantly implementable within some specific experimental scheme. In this paper, we develop a family of such computational tasks (that amount to sampling from suitably prescribed probability distributions). Recently a different approach to similar issues has been described by Aaronson & Arkhipov (2010). More generally there has been increasing interest in physically restricted forms of quantum computing and a study of associated complexity classes (Browne *et al.* 2009; Jordan 2009; Jozsa *et al.* 2009; Shepherd & Bremner 2009; van den Nest 2009).

We consider so-called temporally unstructured quantum computations (also known as IQP or ‘instantaneous’ quantum computation) introduced in Shepherd (2009) and Shepherd & Bremner (2009). Our main result is to demonstrate that if quantum circuits comprising 2-qubit *commuting* gates could be simulated classically (even up to a generous multiplicative error tolerance as described below) then the infinite tower of complexity classes known as the polynomial hierarchy (PH) would collapse to its third level. While not implying that $P = NP$, such a collapse is nevertheless widely regarded as being similarly implausible (Papadimitriou 1994; Arora & Barak 2009). Apart from their tantalizing theoretical simplicity, such circuits of only commuting gates are known to be of significance for super- and semiconductor qubit implementations, where it has recently been shown (Aliferis *et al.* 2009) that they are much simpler to implement fault tolerantly than gates drawn from a fully universal set.

A significant ingredient in our derivations will be the notion of a post-selected quantum computation. Aaronson (2005) has shown that if post-selection is included with universal polynomial time quantum computation then the computational power is boosted from BQP (bounded error quantum polynomial time computation) to the classical class PP (probabilistic polynomial time computation). We will show that, somewhat surprisingly, post-selection boosts the power of the much weaker class of polynomial time IQP computations to PP too.

The notion of classical simulation that applies in our main result is an especially weak one—broadly speaking (see precise definitions below) given a description of a quantum circuit we ask for a classical process that can provide a sample of a probability distribution that approximates the output distribution of the quantum process to a suitable multiplicative accuracy. A more commonly used notion of approximation for probability distributions is that of an *additive* approximation (or approximation in total variation distance),

which often features in the analysis of imperfections in physical implementations. However, there appears to be no useful relationship between additive and multiplicative tolerances, and the proof of our main result remains valid in the context of additive approximations only if the classical simulation is required to be *exact*. The admissibility of multiplicative, in contrast to additive, approximation arises from our use of post-selection; see theorem 3.2 *et seq.*

A very much stronger notion of simulation sometimes used in the literature (which we shall call a strong simulation) is to ask for a classical efficient computation of the *value* of any marginal or total output probability, to exponential precision. Previously it was known (Terhal & DiVincenzo 2004; Fenner *et al.* 2005) that the existence of such strong simulations for some classes of quantum computations would imply the collapse of the PH. Our result contrasts with these works in the striking simplicity of the quantum processes being considered and in the very much weaker requirements in the classical simulation.

2. Preliminary notions

We begin by introducing some definitions and notations needed to give precise statements of our main results.

(a) Computational tasks

Conventionally, a computational task \mathcal{T} is a specified relationship between inputs $w = x_1 \dots x_n$ and outputs $y_1 \dots y_m = \mathcal{T}(w)$ that are taken to be bit strings. The length n of the input string is called the input size. A computational process C with (generally probabilistic) output $C(w)$ on input w is said to compute \mathcal{T} with *bounded error* if there is a constant $0 < \epsilon < 1/2$ such that, for all inputs, $\text{prob}[C(w) = \mathcal{T}(w)] \geq 1 - \epsilon$. C computes \mathcal{T} with *unbounded error* if for all inputs we have $\text{prob}[C(w) = \mathcal{T}(w)] > 1/2$. If the output of \mathcal{T} is always a single bit then \mathcal{T} is called a decision task associated to the subset $\{w : \mathcal{T}(w) = 1\}$ of all bit strings. A subset of bit strings is called a language.

A more general kind of computational task involves merely the sampling of a probability distribution on m -bit strings whose result is not necessarily associated to any desired ‘correct’ outcome $j_1 \dots j_m$ as above. For example, for each n -bit string w , we may have an associated quantum circuit C_w with output probability distribution P_w on m -bit strings, and we may be interested to know how hard it is to sample from P_w by purely classical means, given a description of the circuit C_w .

(b) Uniform families of circuits

We shall use a notion of uniform circuit family that is slightly different from the standard textbook definition, motivated by a desire to make more transparent the contribution of the uniformity condition to the final overall computational process.

In the Turing machine model of computation, a single machine suffices to deal with inputs of all sizes. In contrast in the circuit model, any single circuit has a fixed number of input lines; so, to treat inputs of all sizes it is conventional to introduce the notion of a circuit *family* $\{C_n\} = \{C_1, C_2, \dots\}$ with C_n being a circuit intended to perform the computation for all inputs of size n . In this formalism,

we need to impose an auxiliary *uniformity condition* specifying computational limitations (see below) on how the (descriptions of the) circuits C_n themselves are generated as a function of n . In the absence of any such condition, hard (or even uncomputable) computational results may, by fiat, be hard wired into the varying structure of the circuits C_n with n . In standard treatments, a circuit family $\{C_n\}$ is parametrized by input size n (with C_n being a circuit processing all inputs of size n). For our purposes it will be more convenient to parametrize the circuit family by the inputs $w = x_1 \dots x_n$ themselves, with circuits always acting on a standard input such as $0 \dots 0$ (or $|0\rangle \dots |0\rangle$ for quantum circuits), resulting in circuit families denoted $\{C_w\}$. Thus, for example, in comparison with the standard definition, we could take the circuit C_w to be the circuit C_n prefixed by some NOT gates (depending on w) that initially convert the input $0 \dots 0$ into w . Our formal definition is as follows.

Definition 2.1. A uniform family of circuits (of some specified type) is a mapping $w \rightarrow C_w$ where $w = x_1 \dots x_n$ is a bit string of length n , C_w is a (classical) description of a circuit (of the appropriate type) and the mapping $w \rightarrow C_w$ is computable in classical $\text{poly}(n)$ time. Here, the description C_w includes (i) a specification of a sequence of gates and lines upon which they act, (ii) a specification of the inputs for all lines (often taken to be $0 \dots 0$, respectively, $|0\rangle \dots |0\rangle$ for classical, respectively, quantum circuits), (iii) a specification of which lines constitute the output register, and (iv) a specification of any other registers needed for a circuit of the type being used (e.g. a register of lines initialized to random bit values for randomized computation, or a register of post-selection lines for post-selected computations, as defined later).

Associated to any uniform circuit family, we have a family of probability distributions $\{P_w\}$ (on m -bit strings where m is the size of the output register of C_w), defined by the output of the computational process described by C_w .

Since $w \rightarrow C_w$ is computable in $\text{poly}(n)$ time, each circuit C_w has $\text{poly}(n)$ size and acts on at most $\text{poly}(n)$ lines. One may entertain other uniformity conditions; for example, having $w \rightarrow C_w$ computable in classical log space (as is generally adopted for $n \rightarrow C_n$ in the textbook definition of uniform families). For us the $\text{poly}(n)$ time uniformity condition is adequate, as we are primarily interested in circuits whose computational power is potentially stronger than, or not commensurate with, classical deterministic polynomial time. Our uniformity definition (based on inputs w rather than just input sizes n) then transparently simply prefixes the processing power of the circuits with arbitrary classical deterministic polynomial time computation.

For classical deterministic polynomial time computation in our circuit family definition, the computation can be totally represented within the uniformity stage $w \rightarrow C_w$ and the C_w 's can be taken to be trivial circuits that perform no further computation beyond outputting the obtained answer. Classical *randomized* polynomial time computation is modelled by circuits C_w that have a designated register of lines (disjoint from the input register), which is initialized with random bits for each run of the computation C_w . Such circuits are called classical randomized circuits. The complexity class of decision tasks decided with bounded error (respectively, unbounded error) by uniform families of classical randomized circuits is denoted BPP (bounded error probabilistic polynomial time computation) (respectively, PP). It is well known that BPP is independent of the

value of the constant error tolerance ϵ . For universal polynomial time *quantum* computation the circuits C_w comprise quantum gates, each acting on a constant number of lines. The input is taken to be the standard state $|0\rangle \dots |0\rangle$ and the output is the (probabilistic) result of a computational basis measurement on a designated register of output lines. The class of decision tasks solved with bounded error by such uniform families is denoted BQP. (This definition is easily seen to be equivalent to other standard definitions of BQP such as in Nielsen & Chuang (2000)).

(c) *IQP circuits*

We now come to our notion of quantum computations comprising commuting gates. In Shepherd & Bremner (2009) these have been called IQP ('instantaneous quantum polynomial time') computations since, in quantum physics, such gates may be applied simultaneously.

Definition 2.2. An IQP circuit on n qubit lines is a quantum circuit with the following structure: each gate in the circuit is diagonal in the X basis $\{|0\rangle \pm |1\rangle\}$, the input state is $|0\rangle|0\rangle \dots |0\rangle$ and the output is the result of a computational basis measurement on a specified set of output lines.

In this paper, we will assume that each gate in the description of an IQP circuit C_w is specified by giving its diagonal entries and the lines on which it acts. Thus a poly(n)-sized description implies that any gate acts on at most $O(\log n)$ lines. We note, however, that other inequivalent conventions are possible; for example, in Shepherd & Bremner (2009) gates are specified by a parameter θ and a subset i_1, \dots, i_k of lines, corresponding to the gate $U = \exp(i\theta X_{i_1} \otimes \dots \otimes X_{i_k})$ that may thus act on $O(n)$ lines, although its (potentially exponentially many) diagonal entry phases $\pm\theta$ are all equal up to sign.

It will sometimes be convenient to represent an IQP circuit in terms of gates diagonal in the Z (or computational) basis. In this representation, the inputs and outputs are the same as before but the circuit of gates is required to have the following structure: each qubit line begins and ends with a Hadamard (H) gate, and, in between, every gate is diagonal in the Z basis. This is easily seen to be equivalent to the previous definition (by inserting two H 's on each line between each pair of gates, recalling that $HH = I$, and then absorbing all H 's into conjugation actions on the Z basis diagonal gates, leaving only X basis diagonal gates).

As noted in definition 2.1, any uniform circuit family $\{C_w\}$ associates a probability distribution P_w to each bit string w and we will be especially interested to consider whether this distribution can be sampled (to suitable accuracy) by purely classical means in poly(n) time, given the classical description of the circuit C_w . For this issue it will be significant to note the *number* of output lines, and especially its growth with n .

(d) *Post-selected circuits*

An important theoretical tool in our arguments will be the notion of a *post-selected* (classical or quantum) circuit C . This is a circuit which, in addition to a specified register of output lines \mathcal{O} , has a further (disjoint) specified register of post-selection lines \mathcal{P} . Then instead of sampling measurement results x directly

from the output lines with distribution $\text{prob}[\mathcal{O} = x]$, we consider only those runs of the process for which a measurement on the post-selection lines yields $00\dots 0$, i.e. the output distribution on \mathcal{O} is now taken to be the conditional distribution $\text{prob}[\mathcal{O} = x | \mathcal{P} = 00\dots 0]$. In this construction, we also require the circuit C to have the property that $\text{prob}[\mathcal{P} = 00\dots 0] \neq 0$ so that the conditional probabilities are well defined,

$$\text{prob}[\mathcal{O} = x | \mathcal{P} = 00\dots 0] = \frac{\text{prob}[\mathcal{O} = x \& \mathcal{P} = 00\dots 0]}{\text{prob}[\mathcal{P} = 00\dots 0]}. \quad (2.1)$$

In practical terms a post-selected computation would be implemented by repeatedly running the computation and considering the output register only if the post-selection register is measured to yield $00\dots 0$. Since we place no limit on how small the (non-zero) probability of the latter event may be, the post-selection process may incur an exponential overhead in time, and, similar to the notion of a non-deterministic computation, it is principally of interest as a theoretical tool rather than as a feasible computational resource.

Definition 2.3. A language L is in the class post-IQP (respectively, post-BQP or post-BPP) if and only if there is an error tolerance $0 < \epsilon < 1/2$ and a uniform family $\{C_w\}$ of post-selected IQP (respectively, quantum or randomized classical) circuits with a specified single line output register \mathcal{O}_w (for the L -membership decision problem) and a specified (generally $O(\text{poly}(n))$ -line) post-selection register \mathcal{P}_w such that

- if $w \in L$ then $\text{prob}[\mathcal{O}_w = 1 | \mathcal{P}_w = 00\dots 0] \geq 1 - \epsilon$ and
- if $w \notin L$ then $\text{prob}[\mathcal{O}_w = 0 | \mathcal{P}_w = 00\dots 0] \geq 1 - \epsilon$.

It is pertinent to remark on the ϵ -independence of the classes in definition 2.3 above. The basic bounded error classes BPP and BQP are well known to be independent of the error tolerance $0 < \epsilon < 1/2$. Indeed, the standard method for reducing ϵ is to consider the majority vote answer of multiple runs of the circuit. Similarly, post-BPP and post-BQP are easily seen to be independent of the error tolerance value too. The class post-IQP is in fact also independent of ϵ , as will follow from theorem 3.1 below. However, the class BIQP_ϵ of languages decided with bounded error ϵ by uniform families of IQP circuits (with no post-selection) is not known to be independent of ϵ as it is not evident whether or not the majority vote function can be realized by just (commuting) IQP circuits. Fortunately we will not need to directly use BIQP_ϵ in our arguments.

Post-selected classical computation has been considered in Kuperberg (2009) and Han *et al.* (1997). The class called BPP_{path} that is extensively studied in Han *et al.* (1997) is easily seen to be equal to our class post-BPP.

For quantum computation, the class post-BQP was introduced and studied by Aaronson (2005), who showed that post-BQP equals the classical class PP. Note that, if general quantum or classical circuits are available, it suffices (as in Aaronson (2005)) to use post-selection registers of only a *single* line, since for any register of k lines we may adjoin a circuit that computes some simple function f with $f(x_1 \dots x_k) = 0$ if and only if $x_1 \dots x_k = 00\dots 0$, e.g. the OR of the k bit values suffices. However, if the allowed gates are restricted (as in the case of IQP circuits) it may not be possible to compute any such function using only the allowed resources, and post-selection on multiple lines needs to be entertained, as in our definition above.

(e) *Notions of classical simulation for quantum circuits*

There are various possible notions of classical simulation for quantum circuit families. For any uniform family $\{C_w\}$, let P_w denote the output distribution of C_w and let n denote the length of w .

- (a) We say that a circuit family is *strongly simulatable* if any output probability in P_w and any marginal probability of P_w can be computed to m digits of precision in classical $\text{poly}(n, m)$ time.
- (b) A circuit family is *weakly simulatable* if, given the description of C_w , its output distribution P_w can be sampled by purely classical means in $\text{poly}(n)$ time. Note that strong simulatability implies weak simulatability (Terhal & DiVincenzo 2004)—although the sample space of P_w is exponentially large in n we can sample the distribution in $\text{poly}(n)$ time by successively sampling the bits; the binary distribution used for each successive bit is the conditional distribution, conditioned on the already seen values, and these two conditional probabilities can be computed in $\text{poly}(n)$ time via Bayes' rule, as a quotient of two marginal probabilities of P_w .

Next we have some notions of approximate classical simulation.

- (c) A circuit family is *weakly simulatable with multiplicative error* $c \geq 1$ if there is a family R_w of distributions (on the same sample spaces as P_w) such that R_w can be sampled in classical $\text{poly}(n)$ time and for all x and w we have

$$\frac{1}{c} \text{prob}[P_w = x] \leq \text{prob}[R_w = x] \leq c \text{prob}[P_w = x]. \quad (2.2)$$

- (d) A circuit family is *weakly simulatable within ϵ total variation distance* if there is a family R_w as in (c) above, but with equation (2.2) replaced by the condition

$$\sum_x |\text{prob}[P_w = x] - \text{prob}[R_w = x]| < \epsilon.$$

- (e) A further notion of approximate weak simulation has been formulated in van den Nest (2009): recall first that the Chernoff–Hoeffding bound (see appendix of van den Nest (2009)) implies the following result—if we have a quantum process implementing C_w then by running it poly-many times we can (with probability exponentially close to 1) obtain an estimate \tilde{p} of any output probability p to within polynomial precision, i.e. for any polynomial $f(n)$ we can output \tilde{p} such that $|p - \tilde{p}| < 1/f(n)$. We say that a circuit family is (classically) *weakly simulatable with additive polynomial error* if the same estimates can be obtained from the circuit descriptions C_w by purely classical means in $\text{poly}(n)$ time (and probability exponentially close to 1). Thus, weak simulatability implies weak simulatability with additive polynomial error.

Note that if a uniform circuit family C_w decides a language L with bounded error probability $0 < \epsilon < 1/2$ then the existence of a weak (respectively, strong) simulation for C_w implies that $L \in \text{BPP}$ (respectively, P). Similarly the existence of a weak simulation with additive polynomial error, or with multiplicative error $1 \leq c < 2(1 - \epsilon)$, will also imply that $L \in \text{BPP}$. The latter condition on c serves to guarantee that R_w still decides L with a bounded error $0 < \epsilon' < 1/2$.

3. Main results

(a) *The power of IQP with post-selection*

We begin by examining how the availability of post-selection is able to boost the computational power of various classes of circuits. For this it is convenient to introduce some further notions from complexity theory. If A and B are complexity classes, A^B denotes the class A with an oracle for B (see Papadimitriou (1994) and Arora & Barak (2009) for formal definitions). We may think of A^B as the class of languages decided by the computations subject to the restrictions and acceptance criteria of A but allowing an extra new kind of computational step: we have an oracle or ‘subroutine’ for any desired language L in B that may be queried at any stage in the course of the computation, and each such query counts as a single computational step; that is, bit strings may be generated as intermediate results and presented to the oracle, which, in a single step, returns the information of whether the bit string is in L or not. The PH class (Papadimitriou 1994; Arora & Barak 2009) is defined to be the union of an infinite tower of increasing classes Δ_k , $k = 1, 2, \dots$, in which $\Delta_1 = P$ and $\Delta_{k+1} = P^{N\Delta_k}$. Here, $N\Delta_k$ denotes the non-deterministic class associated to Δ_k , in the same way that NP denotes the non-deterministic class associated to P , i.e. we allow the process to branch at each step into two separate computational paths and deem it to accept its input if and only if at least one path accepts. Further discussion and alternative characterizations of PH may be found in Arora & Barak (2009) and Papadimitriou (1994).

For classical computation it is known (Papadimitriou 1994; Arora & Barak 2009) that BPP is contained in $N\Delta_2$ and also that post-BPP is contained in Δ_3 (Han et al. 1997). Now for any complexity class C we have $P^{(P^C)} = P^C$ (since, in the first expression, any query to a P^C oracle can be replaced by a polynomial time computation with queries to the corresponding oracle for C). Hence we get

$$P^{\text{post-BPP}} \subseteq P^{\Delta_3} = \Delta_3. \quad (3.1)$$

We will use this inclusion below in corollary 3.3.

For the case of quantum computation it is not known whether BQP is contained within PH or not (Aaronson 2009), but, as mentioned above, Aaronson (2005) has shown that $\text{post-BQP} = \text{PP}$. A theorem of Toda (Toda 1991; Papadimitriou 1994; Arora & Barak 2009) asserts that $\text{PH} \subseteq P^{\text{PP}}$ so we get $P^{\text{post-BQP}} \supseteq \text{PH}$. On the other hand, we had $P^{\text{post-BPP}} \subseteq \Delta_3$; so, from an oracle perspective, the power of post-BPP is modest compared with post-BQP or PH.

In view of the above considerations, and recalling that uniform families of IQP circuits are intuitively expected to be far weaker than general quantum computations (and even fail to include many computations in P , such as many elementary arithmetic operations that manifestly depend on the order of operations applied) our next result is perhaps unexpected.

Theorem 3.1. *Post-IQP = post-BQP = PP.*

Proof. Clearly $\text{post-IQP} \subseteq \text{post-BQP}$ and we show the reverse inclusion. Consider an arbitrary uniform quantum circuit family with inputs $|0\rangle \dots |0\rangle$ and with gates drawn from the following universal set: H, Z, CZ and $P = e^{i(\pi/8)Z}$. (For a later purpose we point out here that all these gates are 1- or 2-qubit gates and,




Figure 1. The Hadamard gadget for removal of intermediate H gates. (a) $|\alpha\rangle$ represents a general input state to a gate U within the circuit that is followed by an intermediate H gate. (b) The lower line is a new ancillary qubit line. The original intermediate H gate may then be replaced by a new CZ gate, a post-selection (denoted by $\langle 0|$) and two H gates that are now both at the ends of lines, as allowed in IQP circuit architecture.

apart from H , all gates have diagonal entries that are integer powers of $e^{i\pi/8}$.) If we are allowed to post-select such circuit families then we obtain post-BQP as the class of languages decidable with bounded error. Our strategy is to exhibit a direct reduction from any such post-selected circuit family to a post-selected IQP circuit family whose output conditional probabilities are the same as those of the original family.

Firstly we add in extra H gates to ensure that every line begins and ends with an H gate. This is possible since $H^2 = I$. Next consider in turn each intermediate H gate, i.e. those that do not begin or end a line. For each such gate H_a acting on line a we include an extra qubit line labelled e (for ‘extra’). Consider now the following ‘Hadamard gadget’ illustrated in figure 1. This construction is somewhat akin to gate teleportation and to a basic ingredient underlying measurement-based computation (Nielsen 2006). On lines ae initialized to $|\psi\rangle_a|0\rangle_e$, where $|\psi\rangle$ is any state, we apply the process $|\psi\rangle_a|0\rangle_e \rightarrow H_a CZ_{ae} H_e |\psi\rangle_a|0\rangle_e$ followed by post-selection of outcome 0 on line a . An easy calculation shows that the resulting state on line e is $H|\psi\rangle$. In the original circuit we replace H_a by the Hadamard gadget; here $|\psi\rangle$ represents the circuit’s general input state to H_a and subsequently line e is used as the output line of H_a for further gates in the original circuit. Alternatively we may extend the gadget by a $SWAP_{ae}$ gate and use line a as output. $SWAP$ is not a valid IQP gate so to obtain the final circuit we commute out all $SWAP$ gates to the end of the lines.

In the resulting circuit, the new line e is initialized to $|0\rangle$ and begins and ends with an H gate. Thus, the non-diagonal intermediate H gate has been replaced by a new CZ gate and an additional post-selection. Performing this replacement for every intermediate H gate results in an IQP circuit with some extra post-selections on the new e lines, and with the same output conditional probabilities as originally (now conditioned on the new extra post-selections too). ■

It is interesting to note¹ that there are intuitive analogues in classical complexity theory of the above result—that the availability of post-selection is able to boost the (modest) power of IQP computation to the same level

¹We thank S. Aaronson for drawing our attention to this point.

that it boosts full polynomial-time quantum computation. For example, we may think of 3-CNF Boolean formulae (Arora & Barak 2009) as representing a very simple kind of polynomial-time computation among all classical polynomial-time computations, which is even ‘temporally unstructured’ in the sense that the clauses may be incorporated in any order. Then the Cook–Levin theorem (asserting that 3SAT is NP complete) shows that the availability of non-determinism boosts the power of 3-CNF formulae to the same level that it boosts full polynomial-time classical computation (i.e. to all of NP). Furthermore, the non-deterministic guessing of an accepting path may be thought of as intuitively analogous to the occurrence of a good outcome of the post-selection register measurement.

In the construction used in the above proof, the post-BQP circuit that we start with may, without loss of generality, be assumed to comprise only nearest neighbour 2-qubit gates. Then the SWAP operations introduced by the Hadamard gadgets will at first sight result in a post-IQP circuit that is not truly nearest neighbour. But by simply ‘terminating some of the measurements early’, and ‘creating some ancillas late’, we can avoid line crossings (as is evident from figure 1*b*). The practical outcome of this is that the quantum part of the IQP process resulting from this construction can be rendered, logically speaking, by local interactions on a flat two-dimensional surface (albeit still involving the inefficient resource of post-selection).

(*b*) *Classical simulation of IQP circuits and collapse of PH*

Although IQP circuits have very simple ingredients, we now provide evidence (in corollary 3.3 below) that they, nevertheless, embody computational possibilities that are inaccessible to classical efficient (randomized) computation.

Theorem 3.2. *If the output probability distributions generated by uniform families of IQP circuits could be weakly classically simulated to within multiplicative error $1 \leq c < \sqrt{2}$ then $\text{post-BPP} = \text{PP}$.*

Proof. We will show that, under the stated simulation assumption, any language in post-IQP is in post-BPP and then theorem 3.1 (together with $\text{post-BPP} \subseteq \text{post-BQP}$) will give $\text{post-BPP} = \text{PP}$.

Let $L \in \text{post-IQP}$ be any language decided with bounded error by a uniform family of post-selected IQP circuits C_w with (single line) output registers \mathcal{O}_w and post-selection registers \mathcal{P}_w . Introduce

$$S_w(x) = \frac{\text{prob}[\mathcal{O}_w = x \& \mathcal{P}_w = 0 \dots 0]}{\text{prob}[\mathcal{P}_w = 0 \dots 0]}, \quad (3.2)$$

so the bounded error condition states the following:

$$\left. \begin{array}{l} \text{if } w \in L \text{ then } S_w(1) \geq \frac{1}{2} + \delta \\ \text{if } w \notin L \text{ then } S_w(1) \leq \frac{1}{2} - \delta \end{array} \right\}, \quad (3.3)$$

for some $0 < \delta < 1/2$. Furthermore post-IQP is independent of the level of error; so for any $L \in \text{post-IQP}$ we may assume that equation (3.3) holds for any choice of $0 < \delta < 1/2$, however large. Now let \mathcal{Y}_w denote the full register of lines of C_w , comprising m lines say. If an output measurement on all lines of C_w can be weakly

classically simulated to within multiplicative error c then there is a uniform family of classical randomized circuits \tilde{C}_w with output register $\tilde{\mathcal{Y}}_w$ comprising m lines with

$$\frac{1}{c} \text{prob} [\mathcal{Y}_w = y_1 \dots y_m] \leq \text{prob} [\tilde{\mathcal{Y}}_w = y_1 \dots y_m] \leq c \text{prob} [\mathcal{Y}_w = y_1 \dots y_m]. \quad (3.4)$$

Similarly all marginal distributions for corresponding subregisters of \mathcal{Y}_w and $\tilde{\mathcal{Y}}_w$ satisfy the same inequality. Let $\tilde{\mathcal{O}}_w$ and $\tilde{\mathcal{P}}_w$ denote the registers of \tilde{C}_w corresponding to \mathcal{O}_w and \mathcal{P}_w of C_w , and introduce

$$\tilde{S}_w(x) = \frac{\text{prob} [\tilde{\mathcal{O}}_w = x \& \tilde{\mathcal{P}}_w = 0 \dots 0]}{\text{prob} [\tilde{\mathcal{P}}_w = 0 \dots 0]}. \quad (3.5)$$

Using the inequalities of equation (3.4) (for the registers appearing in equation (3.5)) we get

$$\frac{1}{c^2} S_w(x) \leq \tilde{S}_w(x) \leq c^2 S_w(x). \quad (3.6)$$

Combining this with equation (3.3) we see that the classical uniform family \tilde{C}_w (post-selected on $\tilde{\mathcal{P}}_w$) will decide L with bounded error if $c^2 < 1 + 2\delta$. Since δ can be any value satisfying $\delta < 1/2$ we see that any value of $c < \sqrt{2}$ will suffice to guarantee that $L \in \text{post-BPP}$. ■

It is interesting to point out that our use of a multiplicative approximation (see equation (3.4)) accords well with the quotient structure of the conditional probabilities in equations (3.2) and (3.5), allowing us to derive the bounding relationship equation (3.6) between S_w and \tilde{S}_w . In contrast, use of an additive approximation or approximation to within ϵ total variation distance would be problematic: the denominators of equations (3.2) and (3.5) are required only to be positive, so additive or total variation distance approximations would allow catastrophic divergences of the associated probability quotients.

Corollary 3.3. *If the output probability distributions generated by uniform families of IQP circuits could be weakly classically simulated to within multiplicative error $1 \leq c < \sqrt{2}$ then the PH would collapse to its third level, i.e. $\text{PH} = \Delta_3$.*

Proof. Under the simulation assumption, we may apply theorem 3.2, and Toda's theorem with equation (3.1) gives $\text{PH} \subseteq P^{\text{PP}} = P^{\text{post-BPP}} \subseteq \Delta_3$. ■

From the proof of theorem 3.1 we see that it suffices in theorem 3.2 and corollary 3.3 to require the weak simulatability condition only for a restricted kind of IQP circuit family; namely, those comprising only 1- and 2-qubit gates with diagonal entries being only integer powers of $e^{i\pi/8}$. In a similar vein, one may ask whether the output register may be able to be restricted too, e.g. to having size only $O(\log n)$. Recall that for the class post-IQP, although we have only single-line output registers, the post-selection register may generally have size $O(\text{poly}(n))$ and, in the proof of theorem 3.2, the classical simulation needs to be applicable to IQP circuit families with output registers of the latter size too (as they incorporate the original post-selection registers). Our next result shows that such restriction on the size of the output or post-selection register

is not possible (on the assumption that PH does not collapse), i.e. we see that the computational power of post-selected IQP circuits (with a single line output register) depends crucially on the size of the post-selection register.

Theorem 3.4. *Let P_w be the output probability distributions for any uniform family of IQP circuits in which the output registers have size $O(\log n)$. Then P_w may be sampled (without approximation) by a classical randomized process that runs in time $O(\text{poly}(n))$.*

Proof. Let C_w be any uniform family of IQP circuits with output registers \mathcal{O}_w of size $M = O(\log n)$. Let \mathcal{Y}_w , of size N , denote the complementary register of all non-output lines and let x and y denote generic bit strings of lengths M and N , respectively. We view C_w in its Z -basis diagonal representation: on input $|0\rangle \dots |0\rangle$ the initial Hadamard gates on all lines create an equal superposition and after all Z -diagonal gates of the circuit (and just before the final round of Hadamard gates) the state has the form

$$|\phi\rangle = \frac{1}{\sqrt{2^{M+N}}} \sum_{x,y} e^{if(x,y)} |x, y\rangle. \quad (3.7)$$

The phase function $f(x, y)$ can be computed in classical $\text{poly}(n)$ time by accumulating the relevant diagonal elements of the successive gates. Now the result of further gates and measurements on \mathcal{O}_w is independent of measurements on the disjoint register \mathcal{Y}_w . According to equation (3.7) a measurement of \mathcal{Y}_w will yield a uniformly random bit string of length N . Thus to classically simulate the output of the circuit, we first classically choose a bit string y_0 uniformly at random and consider the state

$$|\phi_{y_0}\rangle = \frac{1}{\sqrt{2^M}} \sum_x e^{if(x,y_0)} |x\rangle.$$

Since $|\phi_{y_0}\rangle$ is a state of only $O(\log n)$ qubits (i.e. $\text{poly}(n)$ dimensions), we can classically strongly simulate the results of further gates and measurements on it, in $\text{poly}(n)$ time by direct calculation, giving overall an exact weak classical simulation of the original circuit family's output. ■

The methods developed in van den Nest (2009) may also be used to readily provide a weak classical simulation up to additive polynomial error for the families in the above theorem.

Shepherd (2010) gives a series of further classical simulation properties of IQP circuits. In particular it is shown there that the distributions P_w in the above theorem are not only exactly weakly simulable, but, even more, they are classically strongly simulable, if all the gates are restricted to have diagonal entries of only integer powers of $e^{i\pi/8}$ (which suffices, as we have noted, to obtain the conclusions of theorems 3.1 and 3.2).

As introduced in Shepherd & Bremner (2009), we may consider a more general notion of an IQP-assisted classical computation than just the single use of the output of a uniform family of IQP circuits. Let \mathcal{IQP} denote an oracle, which, if given a description C of an IQP circuit, will obligingly return (in one computational step) a sample of C 's output distribution. Then we may consider complexity classes such as $\text{BPP}^{\mathcal{IQP}}$, defined as the class of languages decided with

bounded error by a classical probabilistic polynomial-time computation where, in addition to the usual classical steps, the computation may query the oracle with IQP circuit descriptions that have been produced as intermediate results along the way. Since any IQP circuit is a particular kind of quantum circuit, it is easy to see that $\text{BPP}^{\text{IQP}} \subseteq \text{BQP}$, and theorem 3.4 shows that $\text{BPP}^{\text{IQP}[\log n]} = \text{BPP}$, where $\text{BPP}^{\text{IQP}[\log n]}$ denotes that the oracle is queried only with IQP circuits having at most $O(\log n)$ output lines.

4. Some further remarks

It is interesting to note that the methods used to prove our principal results in theorem 3.2 and corollary 3.3 may be applied to other classes of circuits. The only feature of IQP that we needed was the result of theorem 3.1—that post-selection boosts its power to PP. Thus, the evidence of hardness of classical simulation provided by corollary 3.3 would apply to any class of circuits that similarly goes to PP under post-selection. For example, the constructions in Terhal & DiVincenzo (2004) and Fenner *et al.* (2005) (exploiting the notion of gate teleportation (Gottesman & Chuang 1999)) imply that the power of quantum circuits of depth 4 (i.e. three layers of unitary gates followed by a layer of measurements) with post-selection includes BQP and hence also post-BQP = PP, while quantum circuits of depth 3 are known to be always strongly classically simulable. More formally (Fenner *et al.* 2005) introducing the class BQNC^0 of languages decided with bounded error by uniform families of constant depth circuits, we have $\text{post-BQNC}^0 = \text{PP}$ and the conclusion of our corollary 3.3 then applies to QNC^0 (constant depth quantum circuits) replacing IQP.

M.J.B. acknowledges the support of the EU FP7 project COQUIT (contract number 233747), and, in part, of the National Science Foundation under grant no. PHY05-51164 while visiting the KITP. R.J. was supported by the UK EPSRC QIPIRC and the EC network QICS. D.J.S. was funded by CESG. We thank Aram Harrow and Ashley Montanaro for interesting discussions.

References

- Aaronson, S. 2005 Quantum computing, post-selection and probabilistic polynomial-time. *Proc. R. Soc. A* **461**, 3473–3483. (doi:10.1098/rspa.2005.1546)
- Aaronson, S. 2009 BQP and the polynomial hierarchy. (<http://arxiv.org/abs/0910.4698>).
- Aaronson, S. & Arkhipov, A. 2010 New evidence that quantum mechanics is hard to simulate on classical computers. *Talk at QIP2010, Thirteenth International Workshop on Quantum Information Processing, ETH Zurich, Zurich, Switzerland, 18–22 January 2010*.
- Aliferis, P., Brito, F., DiVincenzo, D. P., Preskill, J., Steffen, M. & Terhal, B. M. 2009 Fault-tolerant computing with biased-noise superconducting qubits. *New J. Phys.* **11**, 013061. (<http://arxiv.org/abs/0806.0383>).
- Arora, S. & Barak, B. 2009 *Computational complexity: a modern approach*. Cambridge, UK: Cambridge University Press.
- Browne, D., Kashefi, E. & Perdrix, S. 2009 Computational depth complexity of measurement-based quantum computation. (<http://arxiv.org/abs/0909.4673>).
- Fenner, S., Green, F., Homer, S. & Zhang, Y. 2005 Bounds on the power of constant-depth quantum circuits. In *Proc. 15th Int. Symp. Fundamentals of Computation Theory, Lübeck, Germany, 17–20 August 2005*, pp. 44–55. Lecture Notes in Computer Science. Berlin, Germany: Springer.

- Gottesman, D. & Chuang, I. 1999 Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Lett. Nat.* **402**, 390–393. (doi:10.1038/46503)
- Han, Y., Hemaspaandra, L. & Thierauf, T. 1997 Threshold computation and cryptographic security. *SIAM J. Comput.* **26**, 59–78. (doi:10.1137/S0097539792240467)
- Jordan, S. 2009 Permutational quantum computing. (<http://arxiv.org/abs/0906.2508>).
- Jozsa, R., Kraus, B., Miyake, A. & Watrous, J. 2009 Matchgate and space-bounded quantum computations are equivalent. *Proc. R. Soc. A* **466**, 809–830. (doi:10.1098/rspa.2009.0433)
- Kuperberg, G. 2009 How hard is it to approximate the Jones polynomial? (<http://arxiv.org/abs/0908.0512>).
- Nielsen, M. 2006 Cluster state quantum computation. *Rep. Math. Phys.* **57**, 147–161. (doi:10.1016/S0034-4877(06)80014-5)
- Nielsen, M. & Chuang, I. 2000 *Quantum computation and quantum information*. Cambridge, UK: Cambridge University Press.
- Papadimitriou, C. 1994 *Computational complexity*. Reading, MA: Addison Wesley Publishing Company.
- Shepherd, D. J. 2009 Quantum complexity: restrictions on algorithms and architectures. PhD thesis, University of Bristol, Bristol, UK.
- Shepherd, D. 2010 Binary matroids and quantum probability distributions. (<http://arxiv.org/abs/1005.1744>).
- Shepherd, D. & Bremner, M. J. 2009 Temporally unstructured quantum computation. *Proc. R. Soc. A* **465**, 1413–1439. (doi:10.1098/rspa.2008.0443)
- Terhal, B. & DiVincenzo, D. 2004 Adaptive quantum computation, constant depth quantum circuits and Arthur-Merlin games. *Quant. Inf. Comp.* **4**, 134–145. (<http://arxiv.org/abs/quant-ph/0205133>).
- Toda, S. 1991 PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.* **20**, 865–877. (doi:10.1137/0220053)
- van den Nest, M. 2009 Simulating quantum computers with probabilistic methods. (<http://arxiv.org/abs/0911.1624>).