On The Power of Quantum Fourier Sampling*

(3 page abstract)

Bill Fefferman† Chris Umans‡

1 Introduction

A line of work initiated by Terhal and DiVincenzo [TD02], and Bremner, Jozsa and Shepherd [BJS10] shows that restricted classes of quantum computation can efficiently sample from distributions that cannot be *exactly* sampled classically unless the $\mathbf{PH}$ collapses. The analogous result for decision problems, establishing $\mathbf{BQP} \not\subset \mathbf{BPP}$ unless the $\mathbf{PH}$ collapses, would be a crowning achievement in quantum complexity theory.

Recently there has been a focus on *approximate* sampling problems, with the hope of exhibiting a distribution that can be efficiently sampled quantumly but cannot be even "approximately sampled" classically. This focus is motivated both experimentally, where it seems unreasonable to expect any physical manifestation of the quantum computer *itself* to sample exactly from its distribution, and theoretically, where it has been established by Aaronson [Aar10] that such an approximate hardness result implies the existence of a search problem that can be solved efficiently by a quantum computer but cannot be efficiently solved classically.

Toward this end, Aaronson and Arkhipov [AA13] gave an example of a distribution that can be sampled efficiently by a particular limited form of quantum computation, and which, assuming the validity of two feasible conjectures, cannot be approximately sampled classically (even by a randomized algorithm with a $\mathbf{PH}$ oracle), unless the $\mathbf{PH}$ collapses.

In this work we describe a *general* class of distributions, derived in a natural way from polynomials encoding (presumed) hard problems (of which *Permanent* is but one example) that can be sampled exactly by a quantum computer, but cannot be approximately sampled classically unless the $\mathbf{PH}$ collapses, under variants of the Aaronson-Arkhipov conjectures. Our distributions are described relative to a (presumed) hard polynomial.

This class of polynomials contains the Permanent but also includes, for example, the Hamiltonian Cycle polynomial and many others. We prove our main result by showing that a classical approximate sampler implies an average-case approximation to any such "Efficiently Specifiable" polynomial inside the $\mathbf{PH}$. Since our distribution likely requires the full power of universal quantum computation, while the Aaronson-Arkhipov distribution uses only linear optical quantum computation with noninteracting bosons, why is this result interesting? We can think of at least three reasons:

1. Since the conjectures required in [AA13] have not yet been proven, it seems worthwhile to weaken them as much as possible. We do this in two ways, by weakening both conjectures to apply to any *"Efficiently Specifiable"* polynomial (see Definition 1), and by weakening the so-called Anti-Concentration conjecture so that it need only hold for one distribution in a broad class of distributions.

2. Our construction can be understood without any knowledge of linear optics. While this may be a disadvantage for experimentalists, in our opinion it results in a very clean and simple exposition that may be more immediately accessible to computer scientists.

3. It is extremely common for quantum computations to employ "Quantum Fourier Sampling" in the following way: first apply a classically efficient function to a uniform superposition of inputs, then apply a Quantum Fourier Transform followed by a measurement. Our distributions are obtained in exactly this way, where the classically efficient function is related to a (presumed) hard polynomial. Establishing rigorously a robust sense in which the central primitive of Quantum Fourier Sampling is classically hard seems a worthwhile goal in itself.

2 Efficiently Specifiable Polynomial Sampling on a Quantum Computer

In this section we describe a general class of distributions that can be sampled efficiently on a Quantum Computer. These distributions will form the basis of our main result.

Definition 1 (Efficiently Specifiable Polynomial)**.** *We say a multilinear homogenous n-variate polynomial Q with coefficients in $\{0,1\}$ and m monomials is Efficiently Specifiable if there is an efficiently computable, one-to-one function $h : [m] \to \{0,1\}^n$, with an efficiently computable inverse, and:*

$$Q(X_1, X_2..., X_n) = \sum_{z \in [m]} {X_1}^{h(z)_1} {X_2}^{h(z)_2} ... {X_n}^{h(z)_n}.$$

The class of Efficiently Specifiable polynomials contains the Permanent, the Hamiltonian Cycle polynomial, and other familiar $\#\mathbf{P}$-hard polynomials.

*This is an extended abstract, for the full version see www.math.uchicago.edu/~bill/sampling.pdf

†Supported by NSF CCF-1423544 and BSF grant 2010120.

‡Supported by NSF CCF-1423544 and BSF grant 2010120.

Definition 2 ($\mathcal{D}_Q$)**.** *Suppose Q is an Efficiently Specifiable polynomial with n variables and m monomials. We define distribution $\mathcal{D}_Q$ over binary strings y by:*

$$\Pr_{\mathcal{D}_Q}[y] = \frac{Q((-1)^{y_1}, \ldots, (-1)^{y_n})^2}{2^n m}.$$

Theorem 3 (Quantum Sampling Theorem)**.** *Given a polynomial Q, with n variables, m monomials, that is Efficiently Specifiable via a function h, the distribution $\mathcal{D}_Q$ can be sampled in $poly(n)$ time on a Quantum Computer.*

Proof. We think of our quantum computer as having two registers, the first with $\log m$ qubits and the second with n qubits.

- We start in a uniform superposition over the first register: $\frac{1}{\sqrt{m}} \sum_{z \in [m]} |z\rangle |0..0\rangle$.
- Because h and h^{-1} are efficiently computable we can prepare $\frac{1}{\sqrt{m}} \sum_{z \in [m]} |h(z)\rangle$, by querying h with the first register as input and the second as output, and then h^{-1} with the second as input and the first as output, and then by discarding the first register.
- Apply $H^{\otimes n}$, the Quantum Fourier Transform over $\mathbb{Z}_2^n$, to obtain: $\frac{1}{\sqrt{2^n m}} \sum_{y \in \{0,1\}^n} \sum_{z \in [m]} -1^{\langle y, h(z) \rangle} |y\rangle$

Notice that the squared amplitude of basis vector y in the final state after Step 3 equals the the probability of y under $\mathcal{D}_Q$.[1] □

3 Classical Hardness of Efficiently Specifiable Polynomial Sampling

We are interested in demonstrating the existence of a distribution that can be sampled exactly by a uniform family of quantum circuits, that cannot be sampled approximately classically. Thus the object we wish to rule out is a *Sampler*, defined next, which is a classical algorithm that approximately samples (in total variation distance) from a given class of distributions.

Definition 4 (Sampler)**.** *Let $\{D_n\}_{n>0}$ be a family of distributions where each D_n is distributed over $\mathbb{C}^n$.*

We say S is an (r, ϵ)-Sampler with respect to $\{D_n\}$ if $\|S(0^n, x \sim U_{\{0,1\}^{r(n)}}, 0^{1/\epsilon(n)}) - D_n\| \leqslant \epsilon(n)$ and S runs in (classical) polynomial time.

As in [AA13], our main technical result is that an approximate sampler for $\mathcal{D}_Q$ in the above sense implies an efficient average-case, approximate computation of $|Q|^2$ in the $\mathbf{PH}$, which represents a classical hardness consequence in the case that computing $|Q|^2$ in this fashion is $\#\mathbf{P}$-hard. The proof uses Stockmeyer's Algorithm for approximate counting with an $\mathbf{NP}$ oracle [Sto85]. We say that a computation of $|Q|^2$ is δ-average-case if it succeeds on all but a δ fraction of the inputs x drawn from a specified input distribution, and that it is ϵ-approximate if when it succeeds, the output value is within ϵ of the true value, $|Q(x)|^2$. This error may be *additive* or *multiplicative* (i.e. relative), and we distinguish between these two cases by saying *additive-approximate* or *multiplicative-approximate*.

Theorem 5 (Main)**.** *Given an Efficiently Specifiable polynomial Q with n variables and m monomials, and an $(poly(n), \epsilon\delta)$-Sampler S with respect to $\mathcal{D}_Q$, there is a randomized $(\epsilon m)-$additive approximate $O(\delta)-$average case solution to $|Q|^2$ function with respect to the input distribution $\mathcal{U}_{\{\pm 1\}^n}$, that runs in randomized time $poly(n, 1/\epsilon, 1/\delta)$ with access to an $\mathbf{NP}$ oracle.*

If σ^2 is the variance of the distribution induced by evaluating the polynomial Q with m monomials at a uniformly distributed ± 1-vector (it is easy to calculate $\sigma^2 = m$) then Theorem 5 promises us we can achieve an $\epsilon\sigma^2$-additive approximation to $|Q|^2$, given a classical Sampler for $\mathcal{D}_Q$. The next conjecture asserts, essentially, that the Chebyshev inequality in this setting is tight.

Conjecture 1 (Anti-Concentration Conjecture relative to an n-variate polynomial Q and distribution $\mathcal{D}$ over $\mathbb{C}^n$)**.** *There exists a polynomial p such that for all n and $\delta > 0$,*

$$\Pr_{X \sim \mathcal{D}} \left[|Q(X)|^2 < \frac{\mathrm{Var}\left[Q(X)\right]}{p(n, 1/\delta)} \right] < \delta$$

If this conjecture holds with respect to any *particular* Q and any *particular* input distribution $\mathcal{D}$ considered in this paper (either the uniform distribution on ± 1 vectors discussed in this section, or the binomial distribution considered in the next section), then the additive approximation in Theorem 5 can be replaced with a multiplicative approximation that is more natural when trying to establish $\#\mathbf{P}$-hardness for average-case, approximate, computation of $|Q|^2$.

Theorem 6. *Suppose Conjecture 1 holds relative to an Efficiently Specifiable polynomial Q and an input distribution $\mathcal{D}$, and let X be a random variable distributed according to $\mathcal{D}$. If there is an efficient $\epsilon \mathrm{Var}\left[Q(X)\right]$-additive approximate δ-average case solution to $|Q|^2$ with respect to $\mathcal{D}$, then there is an efficient ϵ'-multiplicative approximate δ'-average case solution to $|Q|^2$ with respect to $\mathcal{D}$, for $\epsilon' = poly(n) \cdot \epsilon$ and $\delta' = 2\delta$. In both cases "efficient" means randomized time $poly(n, 1/\epsilon, 1/\delta)$ with access to an $\mathbf{NP}$ oracle.*

If Conjecture 1 holds with respect to some Efficiently Specifiable polynomial Q and the input distribution of random ± 1 vectors, then establishing the following hardness result, which seems plausible when Q is a $\#\mathbf{P}$-hard function, would be all that is needed to conclude that there is a distribution that can be efficiently quantumly sampled, but that cannot be even approximately sampled classically, unless the $\mathbf{PH}$ collapses.

[1] We also note that by replacing $H^{\otimes n}$ with the Quantum Fourier Transform over $\mathbb{Z}_\ell^n$ and qudits of dimension ℓ, we can sample from a distribution with probabilities proportional to Q evaluated at n-tuples of ℓ-th roots of unity.

Conjecture 2 (Classical Hardness Conjecture)**.** *There exists some Efficiently Specifiable polynomial Q on n variables so that ϵ-multiplicative δ-average case approximation to $|Q|^2$ with respect to the input distribution of random ± 1 vectors cannot be computed in classical randomized time $poly(n, 1/\epsilon, 1/\delta)$ with a $\mathbf{PH}$ oracle.*

Finally, we note that it is interesting to compare Conjecture 1 to the following result of Tao and Vu regarding the Permanent:

Theorem 7 (Tao & Vu [TV08])**.** *For all $\epsilon > 0$ and sufficiently large n,*

$$\Pr_{X \in \{\pm 1\}^{n \times n}} \left[|\mathbf{Permanent}[X]| < \frac{\sqrt{n!}}{n^{\epsilon n}} \right] < \frac{1}{n^{0.1}}$$

This comes quite close to our conjecture for the case of the Permanent polynomial. To our knowledge this bound has not been established for the Gaussian random matrix ensembles considered in [AA13], although as stated there the two distributions should intuitively have similar properties.

4 Distributions Involving Integer Evaluations of Efficiently Specifiable Polynomials

One of the challenges that arises in trying to prove Conjecture 2 is executing a worst-case to average-case reduction for functions, such as the ones considered in this paper, that are not defined over finite fields. To allow for the possibility of mimicking some of the ideas used in such reductions over finite fields (and to obtain the hardness results of [AA13] concerning *exact* average case solutions), it seems useful to consider input distributions beyond those that are uniform ± 1 in each coordinate. In this section we consider input distributions whose support (in each coordinate) is the set of integers in $[-k, k]$, for polynomially bounded k. We do this by means of a reduction: we show a simple way to take an Efficiently Specifiable polynomial with n variables and create another Efficiently Specifiable polynomial with kn variables, in which evaluating this new polynomial at $\{-1, +1\}^{kn}$ is equivalent to evaluation of the old polynomial at $[-k, k]^n$.

Definition 8 (k-valued equivalent polynomial)**.** *For an Efficiently Specifiable polynomial Q with m monomials and an integer $k > 0$, consider the polynomial $Q_k : \{\pm 1\}^{kn} \to \mathbb{R}$ defined by substituting for each variable x_i in Q the sum of k new variables $x_i^{(1)} + x_i^{(2)} + ... + x_i^{(k)}$. We will call Q_k the k-valued equivalent polynomial with respect to Q.*

Note that a uniformly chosen ± 1 assignment to the variables in Q_k induces an assignment to the variables in Q distributed according to the following distribution:

Definition 9 (Distribution $\mathcal{B}(0, k)$)**.** *For k an even integer, we define the distribution $\mathcal{B}(0, k)$ over $[-k, k]$, so that:*

$$\Pr_{\mathcal{B}(0,k)}[y] = \begin{cases} \frac{\binom{k}{(k+y)/2}}{2^k} & \text{if } y \text{ is even} \\ 0 & \text{otherwise} \end{cases}$$

We remark that as k grows $\mathcal{B}(0, k)$ gets closer and closer to the Gaussian distribution with mean 0 and variance 1, which is the distribution considered in [AA13]. In the next two theorems we obtain analogous results to the previous section with respect to the distribution $\mathcal{B}(0, k)$; the proofs are not entirely analogous however, and require some new ideas.

Theorem 10 ($\pm k$-valued Quantum Sampling Theorem)**.** *Given an Efficiently Specifiable polynomial Q with n variables and m monomials, let Q_k be its k-valued equivalent polynomial. We can quantumly sample from the distribution $\mathcal{D}_{Q_k}$ in time $poly(n, k)$.*

Theorem 11 (Classical Consequences Theorem for k-valued sampling)**.** *Let $\mathrm{Var}[Q(X)] = \mathrm{Var}[Q(X_1, X_2, ..., X_n)]$ denote the variance of the distribution induced by Q with assignments drawn from the distribution $\mathcal{B}(0, k)^n$. If there is a $(poly(n), \epsilon\delta)$-Sampler with respect to $\mathcal{D}_{Q_k}$, then there is a randomized $\epsilon \mathrm{Var}[Q(X)]$-additive approximate $O(\delta)$-average case solution to $|Q|^2$ with respect to input distribution $\mathcal{B}(0, k)^n$ that runs in time $poly(n, 1/\epsilon, 1/\delta)$ with access to an $\mathbf{NP}$ oracle.*

As before, Theorem 6 gives us multiplicative approximation if Conjecture 1 holds. As a consequence, under an analogue of Conjecture 2 for the distribution $\mathcal{B}(0, k)$, we would obtain the desired distribution that can be quantumly sampled but not classically approximated. As noted, Conjecture 2 might be easier to prove with respect to the $\mathcal{B}(0, k)^n$ input distribution.

5 Sampling using the "Squashed QFT"

Our final result concerns the input distribution of the previous section, but with k as large as $\exp(n)$. Given an Efficiently Specifiable polynomial Q with n variables, and its k-valued Equivalent Polynomial Q_k, using the prior quantum algorithm of Section 2 we need to invoke the QFT over $\mathbb{Z}_2^{kn}$, which requires $k \leqslant poly(n)$. Can we handle values of k as large as $\exp(n)$? Note that of the 2^{kn} possible $\{\pm 1\}$ assignments to Q_k , there are only k^n *distinct* evaluations. Interestingly, we can take advantage of these massive symmetries, by defining a new unitary operator that can be derived from the standard Quantum Fourier Transform over $\mathbb{Z}_2^n$. We call this the "Squashed QFT". In our paper, we describe the unitary matrix which implements the Squashed QFT, and show how to use it to sample from distributions whose probabilities are proportional to $[-k, k]^n$ evaluations of Efficiently Specifiable polynomials, for $k \leqslant \exp(n)$. Using this construction, and assuming the existence of an efficient quantum circuit for this unitary (which we leave as an open question), we can weaken Conjecture 2:

Conjecture 3 (Weakened Classical Hardness Conjecture)**.** *There exists some Efficiently Specifiable polynomial Q on n variables, and $k \leqslant \exp(n)$ so that ϵ-multiplicative δ-average case approximation to $|Q|^2$ with respect to the input distribution $\mathcal{B}(0, k)^n$ cannot be computed in classical randomized time $poly(n, 1/\epsilon, 1/\delta)$ with a $\mathbf{PH}$ oracle.*

Recall that together with Conjecture 1 this average-case approximate hardness would yield the desired distribution, one that can be efficiently quantumly sampled but not classically approximated.

References

[AA13] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9:143–252, 2013.

[Aar10] Scott Aaronson. The equivalence of sampling and searching. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:128, 2010.

[BJS10] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. 2010.

[Sto85] Larry J. Stockmeyer. On approximation algorithms for #p. *SIAM J. Comput.*, 14(4):849–861, 1985.

[TD02] Barbara M. Terhal and David P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and arthur-merlin games. *CoRR*, quant-ph/0205133, 2002.

[TV08] Terence Tao and Van Vu. On the permanent of random bernoulli matrices. In *Advances in Mathematics*, page 75, 2008.