

Approximate sampling from decoded quantum interferometry via Markov chain Monte Carlo methods [1]

Elies Gil-Fuster,^{1,2,3} Matan Ninio,⁴ Lennart Bittel,¹ Yishai Shimon,⁴
Jens Eisert,^{1,2,5} Stefan Woerner,³ and Almudena Carrera Vázquez³

¹*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*

²*Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany*

³*IBM Research, Rüschlikon 8803, Switzerland*

⁴*IBM Research, Haifa 3498825, Israel*

⁵*Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany*

Combinatorial optimization offers a host of computational problems for which quantum computers may offer an advantage [2]. As such, several attempts have been made in the form of concrete quantum algorithms for optimization [2]. One recent framework, dubbed *Decoded Quantum Interferometry* (DQI) [3], has attracted much attention. Instead of outputting an individual solution, DQI induces a probability distribution biased toward good solutions. DQI is claimed to achieve quantum advantage in a concrete optimization problem called *Optimal Polynomial Intersection* (OPI) [3]. Also, DQI may enable *verifiable* quantum advantage before other algorithms like Shor’s algorithm [4]. Still, it remains challenging to point out the individual quantum ingredients that yield the purported advantage. While certain progress has been made in the front of analytically characterizing the complexity of DQI [5], no final verdicts have been reached. On the empirical front, precious little is known about precisely *how* DQI performs against classical algorithms at different scales. In this work, we shed empirical light on the complexity of DQI. We propose using classical *sampling* algorithms to tackle *optimization* problems, emulating the inner workings of the DQI algorithm.

Results: We employ usual forms of *Markov Chain Monte Carlo* (MCMC) methods to simulate the output state of the DQI algorithm. We numerically test the performance of MCMC methods across sizes, reaching over 100 qubits in the regime where the claims for exponential advantage hold for DQI. We also provide a simplified analytical characterization of the performance of DQI which may be of independent interest.

Our numerical results offer the first glimpses into the complexity of the optimization problems tackled by DQI. We do not challenge the claims of exponential advantage by DQI. Still, we provide evidence that classical algorithms may be able to match the performance of DQI in a reasonable amount of time for a wide range of instance sizes.

Main ideas: We identify a hierarchy of tasks within combinatorial optimization: decision, search, and sampling problems, as depicted in Fig. 1. We carefully analyze the performance of classical *sampling* algorithms on both the *search* and *sampling* version of combinatorial optimization problems (the *decision* version turns out to be trivial in this case). We provide a new characterization of the expected performance of DQI, linking it to usual binomial distributions. We identify a clear obstacle when trying to prove that sampling from the DQI state may be actually classically hard.

Potential impact: Our large-scale numerics serve as direct evidence that classical MCMC algorithms can match the approximation ratio achieved by DQI in practice, for a wide range of instance sizes. We can then use these to discuss the timely differences between sampling and expectation values as different computational primitives. Further, our flexible open-source codebase (in preparation) will enable further benchmarks of other classical sampling algorithms against DQI.

While the optimization problems tackled in this work do not directly resemble usual loss functions in QML, we believe in the fundamental value of our study. Optimization is a core subroutine within larger machine-learning pipelines. We thus expect to gain deep insights for learning by improving our understanding of potential quantum advantages in optimization.




Figure 1. **Hierarchy of tasks in combinatorial optimization.**

I. Combinatorial Optimization and Decoded Quantum Interferometry. We tackle an optimization problem called max-LINSAT. Let $B \in \mathbb{F}_p^{m \times n}$ be a matrix over the cyclic ring $\mathbb{F}_p$, with p being a prime. We say n is the number of variables and m is the number of constraints. For each $i \in [m]$, let $F_i \subseteq \mathbb{F}_p$ be a subset of fixed size $|F_i| = r \leq p$. Then, max-LINSAT is specified as

$$\begin{cases} B_{1,1}x_1 + \dots + B_{1,n}x_n & \in F_1 \\ & \vdots \\ B_{m,1}x_1 + \dots + B_{m,n}x_n & \in F_m, \end{cases} \quad f(x) = \sum_{i=1}^m f_i([Bx]_i) = 2s(x) - m, \quad f_i(a) = \begin{cases} +1 & \text{if } a \in F_i \\ -1 & \text{if } a \notin F_i. \end{cases} \quad (1)$$

The objective function $f(x)$ counts each satisfied constraint as $+1$ and each unsatisfied constraint as -1 . We use $s(x)$ for the number of satisfied constraints. The goal is to maximize $f(x)$: to find an $x \in \mathbb{F}_p^n$ to satisfy as many constraints as possible. This problem is believed to be NP-hard, so we do not expect to find the exact optima efficiently. Instead, we rank different solutions based on their *approximation ratio*: the fraction of satisfied constraints $s(x)/m$.

To tackle this problem, the DQI algorithm prepares the quantum state $|\text{DQI}\rangle = \sum_{x \in \mathbb{F}_p^n} P(f(x))|x\rangle$, where $P(f(x)) = \sum_{k=0}^{\ell} \alpha_k f^k(x)$ is a degree- ℓ polynomial of the objective function. Measuring $|\text{DQI}\rangle$ in the computational basis yields outcome x with probability $\mathcal{P}_{\text{DQI}}(x) = P^2(f(x))$. By carefully choosing the polynomial, DQI induces a distribution over possible solutions which is biased toward higher objective values. The figure of merit is then the expected number of satisfied constraints under this distribution $\langle s \rangle_{\text{DQI}} = \mathbb{E}_{x \sim \mathcal{P}_{\text{DQI}}}[s(x)]$. For a polynomial of degree ℓ , and assuming $r \simeq p/2$, Ref. [3] proves that $\langle s \rangle_{\text{DQI}}/m \simeq 1/2 + \sqrt{\ell/m(1 - \ell/m)}$, referred to as the *semi-circle law*. We do not dwell on how the polynomial is chosen, though it is well understood: the degree ℓ relates to the decoding problem in classical error correction, for the linear code with parity-check matrix $B^\top$; the coefficients α_k can be optimized efficiently to maximize $\langle s \rangle_{\text{DQI}}$. Being precise, preparing $|\text{DQI}\rangle$ requires an efficient decoder that can correct all errors of weight $\leq \ell$. Note that the amplitudes of $|\text{DQI}\rangle$ are efficiently classically computable, and hence also the probabilities $\mathcal{P}_{\text{DQI}}$.

II. Classical Sampling Algorithms. Let $\mathbb{P}$ be a probability distribution for $x \in \mathbb{F}_p^n$ for which we can compute individual probabilities $\mathbb{P}(x)$, and our goal is to produce samples from this distribution. One common approach for this task is Markov Chain Monte Carlo (MCMC) methods. These are iterative processes with stochastic updates which work well heuristically. From the internal state x_t at step t , the next step x_{t+1} is chosen randomly by evaluating $\mathbb{P}$ on a few points around x_t . Different MCMC methods use different update rules, but generally they are all biased towards ensuring $\mathbb{P}(x_{t+1}) > \mathbb{P}(x_t)$. We apply standard MCMC methods directly on the output distribution of DQI, since we noted that we can evaluate $\mathcal{P}_{\text{DQI}}(x)$ efficiently for every $x \in \mathbb{F}_p^n$.

After initial exploration of different methods, we use classical block-Gibbs sampling with block size 3. At step t , we randomly select 3 entries of x_t , and then sample those entries from the exact conditional distribution, conditioned on the values of the other entries of x_t . Consecutive steps thus agree on all entries save for up to 3.




Figure 2. **Optimization-centric metric.** a) For max-XORSAT, the orange line tracks the average score at different times, the star highlights the average number of steps required on average over independent runs (because the performance guarantees of DQI on max-XORSAT only hold on average). b) As the problem size increases, the average runtime τ_{avg} is well-approximated by a power law of degree 5. c) For OPI, the blue line tracks the worst score at different times, the star highlights the maximum number of steps required over the independent runs (because the performance guarantees of DQI on OPI hold in the worst case). d) As the problem size increases, the average runtime τ_{max} is well-approximated by an exponential 1.1^{n_p} . Here n_p is the equivalent number of qubits corresponding to OPI for prime p .

III. Optimization versus Sampling. As in DQI, we employ a sampling algorithm to solve an optimization problem. Our primary goal is to characterize the ability of classical sampling methods to mimic the output distribution of DQI. For this, we establish two different metrics: an *optimization-based* metric and a *sampling-based* metric. The former involves the search problem of finding an individual solution x whose score matches the expected one of DQI: $s(x) \simeq \langle s \rangle_{\text{DQI}}$. The latter involves producing samples consistent with $\mathcal{P}_{\text{DQI}}$. The sketch in Fig. 1 shows that the sampling task is at least as hard as the search task. One reason for this distinction is that, as pointed out in Ref. [5], accurately sampling from $\mathcal{P}_{\text{DQI}}$ may be classically hard, but this does not mean that the search problem is also hard.

IV. Experimental Design. We focus on the settings described in Ref. [3], which we describe briefly:

Irregular-sparse max-XORSAT: This is the simpler case of max-LINSAT where $p = 2$ and $r = 1$: all the variables are boolean, and there is a single acceptance value associated to each constraint $F_i = \{v_i\}$. We refer to $v = (v_i)_{i=1}^m$ as the right-hand-side (rhs) vector. Here, both the coefficient matrix $B \in \mathbb{F}_2^{m \times n}$ and rhs vector v are sampled randomly. The distribution from where B is drawn is specified by a fine-tuned sparsity profile for both the rows and the columns, as well as a fixed ratio $m/n \approx 2$ as we scale n . This profile is chosen to maximize the performance of DQI. For these matrices, the polynomial degree ℓ has to be estimated numerically, which can be done efficiently. Further, the performance guarantee (the semi-circle law) only holds on average over B and v : $\langle s \rangle_{\text{avg}}/m \simeq 1/2 + \sqrt{\ell/m(1 - \ell/m)}$. No claims for exponential speed-ups are made for DQI on this problem: there exist polynomial-time classical optimization algorithms which match the expected performance $\langle s \rangle_{\text{avg}}$.

Optimal Polynomial Intersection: This is the special case of max-LINSAT where $m = p - 1$ and B is a Vandermonde matrix $B_{i,j} = \gamma^{ij}$, for $\gamma \in \mathbb{F}_p^*$ a primitive element. Each of the constraints $\sum_{j=1}^n B_{i,j}x_j$ can be interpreted as the value of a polynomial $Q(y) = \sum_{k=0}^{n-1} Q_k y^k$, where $Q_k = x_k$ and $y = \gamma^i$, hence the name. We consider the parameter regime where both $r/p \simeq 1/2$ and $n/p \simeq 1/2$, as we scale p . For these matrices, the polynomial degree ℓ is known to be $\ell = \lfloor (n - 1)/2 \rfloor$. The search space for this problem is of size p^n , so we introduce the notation n_p as the equivalent number of qubits $p^n \simeq 2^{n_p}$. The performance guarantee here holds worst-case over choices of B (as a function of γ) and the acceptance sets F_i , for $i \in [m]$. An exponential speed-up is claimed for DQI on this problem [3]: the expected performance of DQI is significantly better than the performance of all known polynomial-time classical algorithms.

For both problems, we let MCMC run until success, and we report the number of steps of the Markov chain τ . For the optimization-centric metric, success is defined as finding a single x whose score exceeds the expected DQI score $s(x) \geq \langle s \rangle_{\text{DQI}}$. For the sampling-centric metric, we require the Markov chain to produce 10 *different* good samples (their scores must exceed the expected DQI score). We consider two different sampling algorithms based on the same MCMC, which we dub “restart” (generate each sample independently), and “keep-going” (do not stop iterating after a good sample is found). For max-XORSAT, we tackle the optimization task for up to $n \sim 10^3$, and the sampling task for up to $n \sim 750$. For OPI, we tackle both task for up to $n_p \sim 140$. Our results are presented in Figs. 2 (optimization-centric metric) and 3 (sampling-centric metric).

V. Discussion and outlook. Our experiments reveal that classical sampling algorithms can indeed match the performance of DQI for optimization. Specifically for OPI, the settings where quantum advantage is claimed, we observe an exponential runtime $\tau \propto 1.1^{n_p}$. Still, the base of the exponential is rather small, offering a nuanced discussion on practical quantum advantage with DQI.




Figure 3. **Scaling-centric metric.** a) For max-XORSAT, we report the average runtime to produce different good samples. b) As the problem size increases, both algorithms behave qualitatively similarly, with “keep-going” being faster by a constant multiplicative factor. c) For OPI, we report the average runtime to produce different good samples. d) As the problem size increases, both algorithms behave qualitatively similarly, with “keep-going” being slower by a constant multiplicative factor.

-
- [1] E. Gil-Fuster, M. Ninio, L. Bittel, Y. Shimoni, J. Eisert, S. Woerner, and A. Carrera Vázquez, In preparation (2026).
 - [2] A. Abbas, A. Ambainis, B. Augustino, A. Bärtschi, H. Buhrman, C. Coffrin, G. Cortiana, V. Dunjko, D. J. Egger, B. G. Elmegreen, N. Franco, F. Fratini, B. Fuller, J. Gacon, C. Gonciulea, S. Gribbling, S. Gupta, S. Hadfield, R. Heese, G. Kircher, T. Kleinert, T. Koch, G. Korpas, S. Lenk, J. Marecek, V. Markov, G. Mazzola, S. Mensa, N. Mohseni, G. Nannicini, C. O’Meara, E. P. Tapia, S. Pokutta, M. Proissl, P. Rebentrost, E. Sahin, B. C. B. Symons, S. Tornow, V. Valls, S. Woerner, M. L. Wolf-Bauwens, J. Yard, S. Yarkoni, D. Zechiel, S. Zhuk, and C. Zoufal, *Nature Rev. Phys.* **6**, 718–735 (2024).
 - [3] S. P. Jordan, N. Shutty, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, and R. Babbush, [arXiv:2408.08292 \(2025\)](#).
 - [4] T. Khattar, N. Shutty, C. Gidney, A. Zalcman, N. Yosri, D. Maslov, R. Babbush, and S. P. Jordan, [arXiv: 2510.10967 \(2025\)](#).
 - [5] K. Marwaha, B. Fefferman, A. Gheorghiu, and V. Havlicek, [arXiv: 2509.14443 \(2025\)](#).