

PHYSICS

Easing the Monte Carlo sign problem

Dominik Hangleiter^{1*}, Ingo Roth¹, Daniel Nagaj², Jens Eisert^{1,3}

Quantum Monte Carlo (QMC) methods are the gold standard for studying equilibrium properties of quantum many-body systems. However, in many interesting situations, QMC methods are faced with a sign problem, causing the severe limitation of an exponential increase in the runtime of the QMC algorithm. In this work, we develop a systematic, generally applicable, and practically feasible methodology for easing the sign problem by efficiently computable basis changes and use it to rigorously assess the sign problem. Our framework introduces measures of non-stoquasticity that—as we demonstrate analytically and numerically—at the same time provide a practically relevant and efficiently computable figure of merit for the severity of the sign problem. Complementing this pragmatic mindset, we prove that easing the sign problem in terms of those measures is generally an NP-complete task for nearest-neighbor Hamiltonians and simple basis choices by a reduction to the MAXCUT-problem.

INTRODUCTION

Quantum Monte Carlo (QMC) techniques are central to our understanding of the equilibrium physics of many-body quantum systems. They provide arguably one of the most powerful workhorses for efficiently calculating expectation values of observables in ground and thermal states of various classes of many-body Hamiltonians (1–4). For a Hamiltonian H in dimension D , the idea at the heart of the most prominent variant of QMC is to sample out world lines in a corresponding $(D + 1)$ -dimensional system, where the additional dimension is the (Monte Carlo) time dimension. These world lines correspond to paths through an m -fold expansion of $e^{-\beta H} = (e^{-\beta H/m})^m$, where an entry of $e^{-\beta H/m}$ in a local basis is selected in each step. Each such path is associated with a probability that is proportional to the product of the selected entries. To sample from the resulting distribution, one can construct a suitable Markov chain of paths satisfying detailed balance, which, if gapped, eventually converges to its equilibrium distribution representing the thermal state. Generally speaking, concentration-of-measure phenomena often make such a procedure efficient.

In the classical variant of Monte Carlo, the Hamiltonian is always diagonal, giving rise to positive weights. In QMC, in contrast, positive (in general even complex) off-diagonal matrix elements of H potentially give rise to negative weights of the paths. This leads to what is famously known as the sign problem of QMC, namely, that now one is faced with the task of sampling a quasi-probability distribution (normalized but potentially nonpositive), as opposed to a nonnegative probability distribution. This task can be achieved by introducing a suitable probability distribution that reproduces the desired sampling averages but typically comes at the cost of an exponential increase in the sampling complexity and hence the runtime of the algorithm. For example, in world-line Monte Carlo, one takes the absolute value of the quasi-probability distribution and then computes the average sign, which is given by the expectation value of the signs of the quasi-probabilities with respect to the new distribution. The sign problem is particularly severe for fermionic Hamiltonians, as the particle-exchange anti-symmetry forces their matrix elements to have alternating signs in the standard basis. Naturally, though, it

also appears for bosonic or spin Hamiltonians. The sign problem therefore severely limits our understanding of quantum materials. One can go as far as seeing it to divide strongly correlated systems into easy and intractable cases.

A basic but fundamental insight is that the QMC sign problem is a basis-dependent property (5, 6). For this reason, saying that “a Hamiltonian does or does not exhibit a sign problem” is meaningless without specifying a basis. Because physical quantities of interest are independent of the basis choice, the observation that the sign problem is basis dependent gives immediate hope to actually mitigate the sign problem of QMC by expressing the Hamiltonian in a suitable basis. This is not guaranteed to improve the overall runtime of QMC, as governed not only by the sampling complexity but also by the computational complexity of producing an individual sample. Nonetheless, mitigating the sign problem is widely expected to render QMC efficient in many situations.

In this work, we establish a comprehensive novel framework for assessing, understanding, and optimizing the sign problem computationally, asking the questions: What is the optimal computationally meaningful local basis choice for a QMC simulation of a Hamiltonian problem, can we find it, and how hard is this task in general?

It is known that one can completely cure the sign problem using basis rotations in certain situations. For specific models, sign-problem free bases can be found analytically, involving nonlocal bases, for example, by using so-called auxiliary-field (7), Jordan-Wigner (8), or Majorana (9, 10) transformations. One can also exploit specific known properties of the system such as that the system dimerizes (11–14). These findings motivate the quest for a more broadly applicable systematic search for basis changes that avoid the sign problem, in a way that does not depend on the specific physics of the problem at hand. After all, in a QMC simulation, one wants to learn about the physics of a system in the first place and the optimal basis choice may very well be closely related to that physics.

A useful notion of curing has to restrict the set of allowed basis transformation such that expressing the Hamiltonian in the new basis is still computationally tractable. For example, in its eigenbasis, every Hamiltonian is diagonal and thus sign problem free, but even writing down this basis typically requires an exponential amount of resources. The intrinsic sign problem of a Hamiltonian is thus a property of its equivalence classes under conjugation with some suitable subgroup of the unitary group. The simplest examples of these choices include local Hadamard, Clifford, or unitary transformations. Most

Copyright © 2020
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

¹Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Berlin, Germany. ²RCQI, Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia. ³Helmholtz Center Berlin, Hahn-Meitner-Platz 1, 14109 Berlin, Germany.

*Corresponding author. Email: dominik.hangleiter@fu-berlin.de

generally, one can allow quasi-local circuits that are efficiently computable (6), including not only short circuits and matrix product unitaries (15, 16) but also invertible transformations (17).

A both useful and simple sufficient condition for the absence of a sign problem, independent of the specifics of a simulation, is that the Hamiltonian matrix is stoquastic, i.e., has only nonpositive off-diagonal entries. Stoquasticity provides a useful framework to assess the computational complexity of a systematic approach to curing the sign problem (18). Only recently has the curing problem, to decide whether a stoquastic local basis exists, been shown to be an NP-complete task under on-site unitary transformations for two-local Hamiltonians with additional local fields (19, 20), while it remains efficiently solvable for strictly two-local Hamiltonians (20, 21). But any such approach is faced with the question: Is all hope lost for simulating a Hamiltonian problem via QMC more efficiently even when a stoquastic basis cannot be found in polynomial time?

RESULTS

A pragmatic approach: Easing the sign problem

This leads us to the first part of the initially posed question: What is the optimal computationally meaningful choice of basis? In any Monte Carlo algorithm, computational hardness due to a sign problem is manifested in a super-polynomial increase in its sample complexity as the system size grows. Intuitively speaking, the sample complexity increases because the variance of the Monte Carlo estimator does. In this mindset, finding a QMC algorithm with feasible runtime for Hamiltonians with a sign problem does not require the much stronger task of finding a basis in which the Hamiltonian is fully stoquastic. In many cases, such a basis may not even exist within a given subgroup of the unitaries. Rather, it is often sufficient to merely find a basis in which the Hamiltonian is approximately stoquastic so that the scaling of the variance of the corresponding estimator with the system size is more favorable—ideally polynomial. More pragmatically still, practitioners in QMC are increasingly less worried about small sign problems for which simulations are still feasible for reasonable system sizes using state-of-the-art computing power. This remains true even if the sampling effort may, strictly speaking, diverge exponentially with the system size. Consequently, we argue that practical computational approaches toward the sign problem, rather than focusing on exactly curing it, should target the less ambitious yet practically meaningful task of approximately solving or easing it in the best possible way.

Here, we propose a systematic, generally applicable, and practically feasible methodology for easing the sign problem via basis rotations that allows a meaningful rigorous assessment of this task. An appealing feature of our framework is that it neither requires any a priori knowledge about the physics of a problem nor depends on specifics of a given simulation procedure, in contrast to other known refinements of QMC. At the heart of our approach lies a formulation of the easing problem in terms of a simple, efficiently computable measure of approximate stoquasticity that generically quantifies the sampling complexity.

The sample complexity of a QMC algorithm can be linked to the size of the inverse average sign, which directly bounds the variance of the QMC estimator (18). In an attempt to ease the sign problem of a given Hamiltonian, it is therefore natural to try and improve the average sign. For a few specific models, these improvements have

been achieved by different means: For example, one can exploit known physics to find bases with improved average sign (14, 22) that are often induced by sparse representations (17, 23, 24). For particular observables, one can also exploit clever decompositions of the Monte Carlo estimator into clusters with nonnegative sign (25–31).

However, the sample complexity of computing the average sign via QMC is given by its very value and typically scales exponentially in the system size. Ironically, easing the sign problem by optimizing the average sign is therefore typically infeasible whenever there is a sign problem. One would hence like to quantify the severeness of the sign problem in terms of a quantity that is efficiently computable for physical Hamiltonians—a crucial property to be practically useful in a general approach to easing the sign problem.

Building on the notion of stoquasticity, for a real $D \times D$ Hamiltonian matrix H , we propose the sum of all nonstoquastic matrix entries

$$v_1(H) := D^{-1} \|H_{\rightarrow}\|_{\ell_1} \quad (1)$$

as a natural measure of non-stoquasticity to quantify the sampling complexity of a QMC algorithm in generic instances. Here, as throughout this work, we denote the nonstoquastic part of the Hamiltonian by H_{\rightarrow} , which is defined by $(H_{\rightarrow})_{i,j} = h_{i,j}$ for $h_{i,j} > 0$ and $i \neq j$, and zero otherwise. Moreover, $\|H\|_{\ell_1} = \sum_{i,j} |h_{i,j}|$ is the vector- ℓ_1 -norm.

For local Hamiltonians on bounded-degree graphs such as regular lattices, this measure can be efficiently computed from the nonstoquastic entries of the local terms themselves—for translation-invariant Hamiltonians even with constant effort. However, we can also go beyond that and prove that, for two-local Hamiltonians acting on any graph, the measure v_1 can be efficiently approximated up to any inverse polynomial error (see section S2D in the Supplemental Materials for details). This result renders our measure applicable to problems with long-range and low-degree interactions as they arise, for example, in quantum chemistry.

In principle, one can also conceive of other measures of non-stoquasticity such as the $\ell_1 \rightarrow 1$ -norm or the ℓ_2 -norm of the non-stoquastic part of H . We argue that the ℓ_1 -norm is the most meaningful measure that is agnostic to any particular structure of the Hamiltonian matrix and, therefore, the most versatile measure for a general approach to easing the sign problem. In addition, it acts as a natural regularizer promoting a sparse representation (32) in the spirit of (17, 23, 24).

But how does the non-stoquasticity relate to the sample complexity of a QMC simulation? We find that it is impossible to directly connect a continuous measure of non-stoquasticity to the average sign, which takes on its maximal value at unity and achieves this value for stoquastic Hamiltonians: We can construct exotic examples of highly nonstoquastic Hamiltonians with large positive off-diagonal entries, which also have unit average sign. Conversely, we provide an example of a Hamiltonian with an arbitrarily small non-stoquasticity for which the average sign nearly vanishes.

On the one hand, our examples demonstrate a high sensitivity of the average sign to the Monte Carlo parameters. On the other hand, they also require a malicious interplay between the Hamiltonian matrix entries and highly fine-tuned Monte Carlo parameters. We therefore expect that, in generic situations, the non-stoquasticity measure v_1 meaningfully quantifies the sample complexity of QMC. We give analytical arguments that this is actually the case and numerically find that the average sign of generic two-local Hamiltonians scales exponentially in v_1 (see section S2 in the Supplemental Materials

for details). Thus, we provide evidence that the non-stoquasticity of a local Hamiltonian meaningfully quantifies its sign problem.

Easing in practice

This leads us to the question: Can we practically ease the sign problem of physical Hamiltonians by minimizing non-stoquasticity? To study this second question, we consider translation-invariant nearest-neighbor Hamiltonians in a quasi-one-dimensional geometry (33). Quasi-one-dimensional systems, such as anti-ferromagnetic Heisenberg Hamiltonians on ladder geometries (34, 35), are the simplest nontrivial systems that exhibit a sign problem because they admit the phenomenon of geometric frustration (36). Frustration gives rise to a plethora of phenomena arising in quasi-one-dimensional systems such as the emergence of quantum spin liquids (37, 38) and the interplay of spin-1/2 and spin-1 physics (39). They are also somewhat more realistic descriptions of actual low-dimensional experimental situations than simple one-dimensional chains, serving as a model for small couplings in the transverse direction (35, 40, 41). Therefore, quasi-one-dimensional systems are often seen as a stepping stone toward studying higher dimensions (42), where the sign problem inhibits QMC simulations (43), and thus serve as the perfect playground for a proof of principle.

As a meaningful simple ansatz class, we consider on-site orthogonal transformations $O \in O(d)$ of the type

$$H = \sum_{i=1}^n T_i(h) \mapsto O^{\otimes n} H (O^T)^{\otimes n} \quad (2)$$

for Hamiltonians H acting on n qudits with local dimension d . Here, $T_i(h)$ denotes the translation of a two-local term h to site i . On-site transformations can be handled particularly well as they preserve locality and translation-invariance of local Hamiltonians. In particular, for these transformations, the global non-stoquasticity mea-

sure can be expressed locally in terms of the transformed term $O^{\otimes 2} h (O^T)^{\otimes 2}$ so that the optimization problem has constant complexity in the system size. This constitutes an exponential improvement over approaches that directly optimize the average sign.

To optimize the non-stoquasticity in this setting, we have implemented a geometric optimization method suitable for group manifolds, namely, a conjugate gradient descent algorithm over the orthogonal group $O(d)$ (see section S3 in the Supplemental Materials for details) (44, 45). In Fig. 1A, we show that, generically, the algorithm accurately recovers an on-site stoquastic basis for random Hamiltonians, which are known to admit such a basis a priori. This shows that the heuristic algorithm successfully minimizes the non-stoquasticity and thus serves as a benchmark for its functioning.

We now apply the algorithm to frustrated anti-ferromagnetic Heisenberg Hamiltonians on different ladder geometries (see Fig. 1, B and C). Ladder geometries are interesting not only for the reasons described above but also because, in spite of frustration effects, they often admit sign-problem free QMC methods (11, 13, 14). For both the J_0 - J_1 - J_2 - J_3 model studied in (11) and the frustrated Heisenberg ladder studied in (13, 14), we find a rich optimization landscape in which a relative improvement of the non-stoquasticity by a factor of 2 to 5 can be achieved depending on the region in the phase diagram. Importantly and in spite of those seemingly moderate improvements of non-stoquasticity, we find that the sample complexity of QMC as governed by the inverse average sign is greatly diminished to approximate unity in large regions of the parameter space for the frustrated ladder model (see Fig. 2).

It may well be the case that no stoquastic dimer basis exists, although other variants of QMC do not incur a sign problem for such basis choices: In (11), a stoquastic but nonlocal basis of the J_0 - J_1 - J_2 - J_3 model is identified for values of $J_2 \geq J_0 + J_1$, indicating that more general ansatz classes may well help to further improve the non-stoquasticity. We also observe that first-order optimization algorithms such as the used conjugate gradient method encounter




Fig. 1. We optimize the non-stoquasticity v_1 of translation-invariant, two-local Hamiltonians over on-site orthogonal transformations $O = O^{\otimes n}$ using a conjugate gradient method for manifold optimization (44, 45). (A) Relative non-stoquasticity improvement of random two-local Hamiltonians that are known to admit an on-site stoquastic basis. For each local dimension, 100 instances are drawn and the results are displayed as a box plot according to ([63], 2.16), where whiskers are placed at 1.5 times the interquartile range and circles denote outliers. This serves as a benchmark of our algorithm, which, for almost all instances, accurately recovers a stoquastic on-site basis. (B) Optimized non-stoquasticity of the anti-ferromagnetic J_0 - J_1 - J_2 - J_3 -Heisenberg model relative to the computational basis as a function of $J_2/J, J_3/J$, where $J_0 = J_1 = J$. The algorithm is initialized in a Haar random orthogonal on-site basis. This model is known to admit a nonlocal stoquastic basis for $J_3 \geq J_0 + J_1$ (11). (C) Optimized non-stoquasticity of the anti-ferromagnetic Heisenberg ladder illustrated in the inset with couplings $J_{\parallel}, J_{\perp}, J_x$ relative to the computational basis as a function of J_{\perp}/J_{\parallel} and J_x/J_{\parallel} . We initialized the algorithm at the identity matrix (which was randomly perturbed by a small amount). The phase diagram of the non-stoquasticity qualitatively agrees with the findings of (14), where the stochastic series expansion QMC method was studied. There, it was found that the sign problem can be completely eliminated for a completely frustrated arrangement where $J_x = J_{\parallel}$, while the sign problem remains present for partially frustrated couplings $J_x \neq J_{\parallel}$. However, throughout the parameter regime, the stoquasticity remains nontrivial, which may be due to the fact that the optimization algorithm converges to local minima.




Fig. 2. Improvement of the inverse average sign $(\text{sign})^{-1}$ concomitant with the improvement in non-stoquasticity of Fig. 1C for the frustrated ladder model, as measured by the ratio of its logarithm before optimization compared to that after optimization. We compute the average sign via exact diagonalization for a ladder of 2×4 -sites, $m = 100$ Monte Carlo steps, and inverse temperature $\beta = 1$.

obstacles due to the rugged non-stoquasticity landscape. Intuitively, this landscape is governed by the combinatorial increase of possible assignments of signs to the Hamiltonian matrix elements.

The findings of our proof-of-principle study are twofold: On the one hand, they show that one can efficiently optimize the non-stoquasticity for translation-invariant problems that admit a stoquastic basis lying within the ansatz orbit. They also further substantiate the claim that optimizing non-stoquasticity typically eases the sign problem and dampens the increase in sampling complexity. In addition, they indicate that more general ansatz classes such as quasi-local circuits yield the promise to further reduce the non-stoquasticity of ladder models. We therefore expect that optimizing non-stoquasticity is a feasible and promising means to reduce the sign problem for many different systems, including two-dimensional lattice systems, by exploiting the flexibility offered by larger ansatz classes within our framework. On the other hand, already in our small study, we encountered obstacles preventing efficient optimization of the non-stoquasticity in the guise of a complicated and rugged optimization landscape.

The computational complexity of SignEasing

Fundamentally, our findings thus raise the third question: How far can an approach to easing the sign problem using optimization over local bases carry in principle? In our main complexity-theoretic result, we systematically study the fundamental limits of minimizing non-stoquasticity as a means to ease the sign problem. To do so, we complement the pragmatic mindset of this work with the rigorous machinery of computational complexity theory, asking the question: What is the computational complexity of optimally easing the sign problem? To formalize this question, we introduce the corresponding decision problem:

Definition 1 (SignEasing). Given an n -qubit Hamiltonian H , constants $B > A \geq 0$ with $B - A \geq 1/\text{poly}(n)$, and a set of allowed unitary transformations \mathcal{U} , decide which of the following is the case

$$\text{YES: } \exists U \in \mathcal{U}: v_1(UHU^\dagger) \leq A, \text{ or} \tag{3}$$

$$\text{NO: } \forall U \in \mathcal{U}: v_1(UHU^\dagger) \geq B \tag{4}$$

Table 1. Analogy between satisfiability problems and curing/easing the sign problem. The satisfiability equivalent of curing the sign problem is to decide whether a given sentence is satisfiable, while the equivalent of easing is to find the minimal number of clauses that are violated by a sentence. Similarly, results on the computational complexity of curing and easing the non-stoquasticity of a local Hamiltonian H are in one-to-one correspondence with the hardness of satisfiability problems.

Satisfiability	Stoquasticity	Complexity	References
3SAT	Curing 2 + 1-local H	NP-complete	(19, 20)
2SAT	Curing strictly 2-local H	In P	(20, 21)
MAX2SAT	Easing strictly 2-local H	NP-complete	Here

We derive the computational complexity of the sign easing problem in simple settings, namely, for two-local Hamiltonians, allowing on-site orthogonal Clifford operations and on-site general orthogonal transformations. We prove that under both classes of transformations, SignEasing is NP-complete. Intriguingly, this holds true even in cases in which the curing problem can be decided efficiently, namely, for strictly two-local XYZ Hamiltonians of the type considered in (20, 21).

Theorem 2 (Complexity of SignEasing). SignEasing is NP-complete for 2-local (XYZ) Hamiltonians under

- (i) on-site orthogonal Clifford transformations and
- (ii) on-site general orthogonal transformations.

From a practical perspective, our results pose limitations on the worst-case runtime of algorithms designed to find optimal QMC bases for the physically relevant case of two-local Hamiltonians. From a complexity-theoretic perspective, they manifest a sign-problem variant of the dichotomy between the efficiently solvable 2SAT-problem to decide whether there exists a satisfying assignment for a two-local sentence, and the NP-complete MAX2SAT-problem asking what is the least possible number of broken clauses. They thus complete the picture drawn by (19–21) regarding the connection between satisfiability problems and the problems of curing and easing the sign problem on arbitrary graphs, a state of affairs that we illustrate in Table 1. It is natural to ask the question how far this connection extends and what we can learn from it about efficiently solvable instances. For example, one may ask whether results about the hard regions of 3SAT and MAX2SAT carry over to the problems of curing and easing the sign problem.

MATERIALS AND METHODS

We prove Theorems 2i and 2ii as Theorems 8 and 9 in the Supplementary Materials. The essential idea of our proof, sketched below and illustrated in Fig. 3, is to design a corresponding Hamiltonian such that if the sign problem could be optimally eased for this Hamiltonian under the respective ansatz class, then one could also find the ground-state energy of the original anti-ferromagnetic Ising Hamiltonian, a task that is NP-hard to begin with. It is straightforward to prove versions of Theorem 2 for any ℓ_p -norm of the nonstoquastic part of H , with finite p as a measure of non-stoquasticity. Our result is therefore independent of the particular choice of (ℓ_p) non-stoquasticity measure.




Fig. 3. Constructing a Hamiltonian whose sign problem is NP-hard to ease under orthogonal on-site transformations. (A) To prove NP-completeness of SignEasing, we reduced it to the MAXCUT-problem, which asks for the ground-state energy of an anti-ferromagnetic Ising Hamiltonian H on a graph G . (B) In our encoding, we map H to a Hamiltonian H' , in which all ZZ interactions are replaced by XX interactions and translate the spin configurations $(s_1, \dots, s_n) \in \{0,1\}^n$ of the anti-ferromagnetic Ising model to on-site transformations $Z_1^{s_1} \dots Z_n^{s_n}$. To achieve this restriction, we penalize all other transformations by adding an ancilla qubit $\xi_{i,j}$ for every edge (i,j) of G and adding the interaction term $C(Z_i Z_j - Z_i Z_{\xi_{i,j}} - Z_j Z_{\xi_{i,j}})$ with a suitably chosen constant $C > 0$. We obtain that $v_1(H')$ can be eased below a certain value if and only if the ground-state energy of H is below that value to begin with, thus establishing the reduction.

Proof sketch

SignEasing for arbitrary two-local Hamiltonians is contained in NP—given a basis transformation, we can approximate the measure of non-stoquasticity from the transformed local terms up to any inverse polynomial error and hence verify the YES-case (Eq. 3) (see Theorem 6 in the Supplementary Materials).

The key idea of the harder direction of the proof is to encode the promise version of the MAXCUT-problem into the SignEasing problem. An instance of MAXCUT is given by a graph $G = (V, E)$, and the problem is to decide whether the ground-state energy of the anti-ferromagnetic Ising Hamiltonian

$$H = \sum_{(i,j) \in E} Z_i Z_j \quad (5)$$

is below a constant A or above B . Here, Z_i is the Pauli-Z operator acting on site i . We now define a Hamiltonian H' in which we replace every $Z_i Z_j$ interaction of H by an $X_i X_j$ interaction, as we illustrate in Fig. 3. To understand our embedding, suppose that we perform basis changes only by applying Z or $\mathbb{1}$ at every site. In this case, a Hamiltonian term can be made stoquastic if and only if $X_i X_j \mapsto -X_i X_j$, which is achieved by a transformation $Z^{s_i} Z^{s_j}$ with $(s_i, s_j) = (0,1) \vee (1,0)$. A term remains stoquastic for $(s_i, s_j) = (1,1) \vee (0,0)$. This provides a direct mapping between spin configurations $(1,0)$ and $(0,1)$, which do not contribute to the ground-state energy of the anti-ferromagnetic Ising model and transformations that make local terms in H' stoquastic and thus decrease the non-stoquasticity.

To prove the theorem for arbitrary on-site Clifford and orthogonal transformations, we introduce an additional qubit $\xi_{i,j}$ for every edge (i,j) and add interaction terms $C(Z_i Z_j - Z_i Z_{\xi_{i,j}} - Z_j Z_{\xi_{i,j}})$ to H' with constant $C = 2 \deg(G)$, where $\deg(G)$ is the degree of the interaction graph G (see Fig. 3B). These terms penalize all other transformations such that the optimal non-stoquasticity of H' is always achieved for transformations of the form $Z_1^{s_1} \dots Z_n^{s_n}$ with $(s_1, \dots, s_n) \in \{0,1\}^n$. For example, suppose that we apply Hadamard transformations to all sites $i, j, \xi_{i,j}$, then the ZZ interactions and XX interactions change roles so that the non-stoquasticity cannot be decreased by such a transformation. Showing this for all possible transformations constitutes the main technical part of the proof.

Because MAXCUT is a variant of the MAX2SAT-problem, our results not only manifest but also crucially use the 2SAT-MAX2SAT dichotomy. Notice that because the MAXCUT-problem is NP-hard already on subgraphs of the double-layered square lattice (46),

which has degree six, hard instances of the sign-easing problem occur already for low-dimensional lattices with small (constant) interaction strength.

In our complexity-theoretic analysis, we have focused on the computational complexity of easing the sign problem as the size of an arbitrary input graph is scaled up, in the same mindset as (18–21). We expect, however, that the complexity of SignEasing scales similarly in the size of the lattice unit cell and the local dimension of translation-invariant systems such as those discussed above.

DISCUSSION

Summary

Our work introduces the sign easing methodology as a systematic novel paradigm useful for assessing and understanding the sign problem of QMC simulations. We ask and answer three central questions using complementary methods from theoretical and applied computer science as well as from physics. First, we define a measure of non-stoquasticity suitable for easing the sign problem and extensively discussed its relation to the average sign. Second, we demonstrate that one can feasibly optimize this measure over local bases in simple settings by applying geometric optimization methods. Last, we establish the computational complexity of sign easing in a broader but still simple setting. In this way, our work not only identifies a means of easing the sign problem and demonstrates its feasibility and potential but also shows up its fundamental limitations in terms of computational complexity. Even more so, we are confident that the framework of our work provides both valuable guidance and the practical means for future research on systematically easing the sign problem of Hamiltonians that are particularly interesting and relevant in condensed matter and material science applications.

Outlook

As a first general and systematic attempt to easing the sign problem, we have restricted the focus of this work in several ways. Hence, a number of questions, generalizing our results in different directions, are left open.

First, we have restricted our discussion to the prominent world-line Monte Carlo method to maintain clarity throughout the article. We are confident, however, that our results find immediate application for other Monte Carlo methods such as stochastic series expansion Monte Carlo and determinantal Monte Carlo (36, 47) as well as diffusion Monte Carlo techniques such as full-configuration interaction

Monte Carlo (48). Similar sign problems involving the sampling from quasi-probability distributions also appear in different contexts, for example, in approaches to the classical simulation of quantum circuits (49–51) or high-energy physics (52). In these contexts, too, the problem of finding better bases in which to perform the sampling appears. While the framework developed in this work uses the specific features of QMC, the general idea and mindset behind it applies to all basis-dependent sign problems. Our work thus paves the way toward easing sign problems in a plethora of contexts.

Second, we have only considered real-valued Hamiltonians and transformations that preserve this property. For general complex-valued Hamiltonians, the sign problem takes the form of a complex phase problem. A natural follow-up of our work is to explore how our results on easing the sign problem generalize to the complex phase problem.

Third, we have put an emphasis on the conjugation of Hamiltonians under on-site Clifford and orthogonal circuits. In principle, one may also allow arbitrary quasi-local circuits, as long as the conjugation can be efficiently computed, albeit of exponentially increasing effort with the support of the involved unitaries. This leads to the interesting insight that within the trivial phase of matter, one can always remove the sign problem: One has to conjugate the Hamiltonian with the quasi-local unitary that brings a given Hamiltonian into an on-site form of a fixed-point Hamiltonian. For given Hamiltonians, this may be impractical, of course. In this sense, one can identify trivial quantum phases of matter as efficiently computable phases of matter, an intriguing state of affairs from a conceptual perspective. Conversely, for topologically ordered systems, there may be topological obstructions to curing the sign problem by any quasi-local circuit (6, 53), giving rise to an entire phase of matter that exhibits an intrinsic sign problem. For example, the fixed-point Hamiltonians of the most general class of nonchiral topologically ordered systems, the Levin-Wen models (54), are associated with 12-local Hamiltonians, many of which are expected to not be curable from their sign problem. This insight further motivates to study the sign easing problem for efficiently computable subgroups of local unitaries from a perspective of topological phases of matter.

Our work also opens up several paths for future research. The immediate and practically most relevant direction is, of course, to find the best possible way of minimizing the non-stoquasticity and to explore how well the sign problem can be eased in systems that are not yet amenable to QMC. We have already introduced a flexible optimization approach, which can be straightforwardly applied to a wide range of translation-invariant systems and ansatz classes in any dimensionality. In this respect, it will be interesting to compare possible ways of optimizing the sign problem via different measures (55) and optimization algorithms (56) in various systems (57).

Furthermore, in our hardness proof, we have shown that the easing problem is intricately related to satisfiability problems. Building on this connection, an exciting direction of research is to combine highly efficient SAT-solvers that are capable of exploring combinatorically large sets, with manifold optimization techniques that are able to handle rich geometrical structures, in the spirit of recent work (58). While our hardness result shows up fundamental limitations of SignEasing in the general case, it thus also opens the door to potentially solve the sign easing problem in relevant instances by applying methods well known in computer science to relaxed versions of the easing problem. One may thus hope that for large classes of relevant instances minimizing the non-stoquasticity is actually tractable.

A question closely related to the sign easing problem is the following: How hard is it to find the ground-state energy of a stoquastic Hamiltonian—a subproblem of the so-called local Hamiltonian problem. The computational complexity of this stoquastic local Hamiltonian problem poses fundamental limitations on the classical simulatability of Hamiltonians, which do not suffer from a sign problem and are therefore amenable to QMC simulations. It has been shown that the two-local stoquastic Hamiltonian problem is complete for the class StoqMA (59, 60), a class intermediate between AM and MA that also functions as a genuinely intermediate class in the complexity classification of local Hamiltonian problems (61), even when extending to the full low-energy spectrum (62). The results of (59) also imply that we cannot expect to efficiently find a stoquastic local basis for arbitrary local Hamiltonians unless the unlikely complexity-theoretic equality $AM = QMA$ holds.

For efficiently curable Hamiltonians, the local Hamiltonian problem is reduced to a stoquastic local Hamiltonian problem. Conversely, both the easing problem and the stoquastic local Hamiltonian problem contribute to the hardness of a QMC procedure. For a given Hamiltonian, QMC may thus be computationally intractable for two reasons: it is hard to find a basis in which the Hamiltonian is stoquastic, or cooling to its ground state is computationally hard in its own right. In a QMC algorithm, the latter hardness is manifested as a Markov chain Monte Carlo algorithm not converging in polynomial time. This may be the case even for classical models such as Ising spin glasses (46).

An important open question is how the hardness of easing the sign problem and the hardness of sampling from the estimator distribution are related in specific cases. For example, when improving the average sign, the hardness of a problem that was manifest in an increased sample complexity of the Monte Carlo estimator might be “transferred” to the hardness of sampling from the resulting distribution. On the other hand, there might be instances in which the only obstacle in the way of an efficient simulation is to find a certain basis in which the corresponding Hamiltonian has a large average sign, but, given that basis, QMC runs efficiently.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/6/33/eabb8341/DC1>

REFERENCES AND NOTES

1. J. E. Hirsch, R. L. Sugar, D. J. Scalapino, R. Blankenbecler, Monte Carlo simulations of one-dimensional fermion systems. *Phys. Rev. B* **26**, 5033–5055 (1982).
2. M. Troyer, F. Alet, S. Trebst, S. Wessel, Non-local updates for quantum Monte Carlo simulations. *AIP Conf. Proc.* **690**, 156–169 (2003).
3. L. Pollet, Recent developments in quantum Monte-Carlo simulations with applications for cold gases. *Rep. Prog. Phys.* **75**, 094501 (2012).
4. S. Trotzky, L. Pollet, F. Gerbier, U. Schnorrberger, I. Bloch, N. V. Prokof'ev, B. Svistunov, M. Troyer, Suppression of the critical temperature for superfluidity near the Mott transition. *Nat. Phys.* **6**, 998–1004 (2010).
5. N. Hatano, M. Suzuki, Representation basis in quantum Monte Carlo calculations and the negative-sign problem. *Phys. Lett. A* **163**, 246–249 (1992).
6. M. B. Hastings, How quantum are non-negative wavefunctions? *J. Math. Phys.* **57**, 015210 (2015).
7. C. Wu, J.-P. Hu, S.-C. Zhang, Exact SO(5) symmetry in the spin-3/2 fermionic system. *Phys. Rev. Lett.* **91**, 186402 (2003).
8. K. Okunishi, K. Harada, Symmetry-protected topological order and negative-sign problem for SO(N) bilinear-biquadratic chains. *Phys. Rev. B* **89**, 134422 (2014).
9. Z.-X. Li, Y.-F. Jiang, H. Yao, Solving the fermion sign problem in quantum Monte Carlo simulations by Majorana representation. *Phys. Rev. B* **91**, 241117 (2015).

10. Z.-X. Li, Y.-F. Jiang, H. Yao, Majorana-time-reversal symmetries: A fundamental principle for sign-problem-free quantum Monte Carlo simulations. *Phys. Rev. Lett.* **117**, 267002 (2016).
11. T. Nakamura, Vanishing of the negative-sign problem of quantum Monte Carlo simulations in one-dimensional frustrated spin systems. *Phys. Rev. B* **57**, R3197–R3200 (1998).
12. F. Alet, K. Damle, S. Pujari, Sign-problem-free Monte Carlo simulation of certain frustrated quantum magnets. *Phys. Rev. Lett.* **117**, 197203 (2016).
13. A. Honecker, S. Wessel, R. Kerkdyk, T. Pruschke, F. Mila, B. Normand, Thermodynamic properties of highly frustrated quantum spin ladders: Influence of many-particle bound states. *Phys. Rev. B* **93**, 054408 (2016).
14. S. Wessel, B. Normand, F. Mila, A. Honecker, Efficient quantum Monte Carlo simulations of highly frustrated magnets: The frustrated spin-1/2 ladder. *SciPost Phys.* **3**, 005 (2017).
15. J. I. Cirac, D. Perez-Garcia, N. Schuch, F. Verstraete, Matrix product unitaries: Structure, symmetries, and topological invariants. *J. Stat. Mech.* **2017**, 083105 (2017).
16. M. B. Şahinoğlu, S. K. Shukla, F. Bi, X. Chen, Matrix product representation of locality preserving unitaries. *Phys. Rev. B* **98**, 245122 (2018).
17. W. Dobrautz, H. Luo, A. Alavi, Compact numerical solutions to the two-dimensional repulsive Hubbard model obtained via nonunitary similarity transformations. *Phys. Rev. B* **99**, 075119 (2019).
18. M. Troyer, U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations. *Phys. Rev. Lett.* **94**, 170201 (2005).
19. M. Marvian, D. A. Lidar, I. Hen, On the computational complexity of curing non-stoquastic Hamiltonians. *Nat. Commun.* **10**, 1571 (2019).
20. J. Klassen, M. Marvian, S. Piddock, M. Ioannou, I. Hen, B. Terhal, Hardness and ease of curing the sign problem for two-local qubit Hamiltonians. arXiv:1906.08800.3000 [quant-ph] (20 June 2019).
21. J. Klassen, B. M. Terhal, Two-local qubit Hamiltonians: When are they stoquastic? *Quantum* **3**, 139 (2019).
22. H. Shinaoka, Y. Nomura, S. Biermann, M. Troyer, P. Werner, Negative sign problem in continuous-time quantum Monte Carlo: Optimal choice of single-particle basis for impurity problems. *Phys. Rev. B* **92**, 195126 (2015).
23. J. R. McClean, A. Aspuru-Guzik, Clock quantum Monte Carlo: An imaginary-time method for real-time quantum dynamics. *Phys. Rev. A* **91**, 012311 (2015).
24. R. E. Thomas, Q. Sun, A. Alavi, G. H. Booth, Stochastic multiconfigurational self-consistent field theory. *J. Chem. Theory Comput.* **11**, 5316–5325 (2015).
25. W. Bietenholz, A. Pochinsky, U.-J. Wiese, Meron-cluster simulation of the θ vacuum in the 2D O(3) model. *Phys. Rev. Lett.* **75**, 4524–4527 (1995).
26. S. Chandrasekharan, U.-J. Wiese, Meron-cluster solution of fermion sign problems. *Phys. Rev. Lett.* **83**, 3116–3119 (1999).
27. P. Henelius, A. W. Sandvik, Sign problem in Monte Carlo simulations of frustrated quantum spin systems. *Phys. Rev. B* **62**, 1102–1113 (2000).
28. M. Nyfeler, F.-J. Jiang, F. Kämpfer, U.-J. Wiese, Nested cluster algorithm for frustrated quantum antiferromagnets. *Phys. Rev. Lett.* **100**, 247206 (2008).
29. E. Huffman, S. Chandrasekharan, Solution to sign problems in models of interacting fermions and quantum spins. *Phys. Rev. E* **94**, 043311 (2016).
30. C. T. Hann, E. Huffman, S. Chandrasekharan, Solution to the sign problem in a frustrated quantum impurity model. *Ann. Phys.* **376**, 63–75 (2017).
31. I. Hen, Resolution of the sign problem for a frustrated triplet of spins. *Phys. Rev. E* **99**, 033306 (2019).
32. S. Foucart, H. Rauhut, A mathematical introduction to compressive sensing, in *Applied and Numerical Harmonic Analysis* (Springer, 2013).
33. H.-J. Mikeska, A. K. Kolezhuk, One-dimensional magnetism, in *Quantum Magnetism*, U. Schollwöck, J. Richter, D. J. J. Farnell, R. F. Bishop, Eds. (Lecture Notes in Physics, Springer Berlin Heidelberg, 2004), pp. 1–83.
34. E. Dagotto, T. M. Rice, Surprises on the way from one- to two-dimensional quantum magnets: The ladder materials. *Science* **271**, 618–623 (1996).
35. M. Takano, Spin ladder compounds. *Phys. C* **263**, 468–474 (1996).
36. A. W. Sandvik, Computational studies of quantum spin systems. *AIP Conf. Proc.* **1297**, 135–338 (2010).
37. T. Meng, T. Neupert, M. Greiter, R. Thomale, Coupled-wire construction of chiral spin liquids. *Phys. Rev. B* **91**, 241106 (2015).
38. P.-H. Huang, J.-H. Chen, A. E. Feiguin, C. Chamon, C. Mudry, Coupled spin-1/2 ladders as microscopic models for non-Abelian chiral spin liquids. *Phys. Rev. B* **95**, 144413 (2017).
39. A. Nietner, C. Krumnow, E. J. Bergholtz, J. Eisert, Composite symmetry protected topological order and effective models. *Phys. Rev. B* **96**, 235138 (2017).
40. Y. Yoshida, H. Ito, M. Maesato, Y. Shimizu, H. Hayama, T. Hiramatsu, Y. Nakamura, H. Kishida, T. Koretsune, C. Hotta, G. Saito, Spin-disordered quantum phases in a quasi-one-dimensional triangular lattice. *Nat. Phys.* **11**, 679–683 (2015).
41. K. T. Lai, M. Vallдор, Coexistence of spin ordering on ladders and spin dimer formation in a new-structure-type compound $\text{Sr}_2\text{Co}_5\text{S}_2\text{O}_3$. *Sci. Rep.* **7**, 43767 (2017).
42. V. I. Iglovikov, E. Khatami, R. T. Scalettar, Geometry dependence of the sign problem in quantum Monte Carlo simulations. *Phys. Rev. B* **92**, 045110 (2015).
43. J. Carrasquilla, Z. Hao, R. G. Melko, A two-dimensional spin liquid in quantum kagome ice. *Nat. Commun.* **6**, 7421 (2015).
44. T. Abrudan, J. Eriksson, V. Koivunen, Conjugate gradient algorithm for optimization under unitary matrix constraint. *Signal Process.* **89**, 1704–1714 (2009).
45. D. Hangleiter, I. Roth, Gitlab repository (2019); <https://gitlab.com/ingo.roth/signease>.
46. F. Barahona, On the computational complexity of Ising spin glass models. *J. Phys. A* **15**, 3241–3253 (1982).
47. D. P. Landau, K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge Univ. Press, 2000).
48. G. H. Booth, A. J. W. Thom, A. Alavi, Fermion Monte Carlo without fixed nodes: A game of life, death, and annihilation in Slater determinant space. *J. Chem. Phys.* **131**, 054106 (2009).
49. C. M. Dawson, H. L. Haselgrove, A. P. Hines, D. Mortimer, M. A. Nielsen, T. J. Osborne, Quantum computing and polynomial equations over the finite field \mathbb{Z}_2 . *Quant. Inf. Comp.* **5**, 102–112 (2005).
50. S. P. Jordan, D. Gosset, P. J. Love, Quantum-Merlin-Arthur-complete problems for stoquastic Hamiltonians and Markov matrices. *Phys. Rev. A* **81**, 032331 (2010).
51. H. Pashayan, J. J. Wallman, S. D. Bartlett, Estimating outcome probabilities of quantum circuits using quasiprobabilities. *Phys. Rev. Lett.* **115**, 070501 (2015).
52. K. N. Anagnostopoulos, J. Nishimura, New approach to the complex-action problem and its application to a nonperturbative study of superstring theory. *Phys. Rev. D* **66**, 106008 (2002).
53. Z. Ringel, D. L. Kovrizhin, Quantized gravitational responses, the sign problem, and quantum complexity. *Sci. Adv.* **3**, e1701758 (2017).
54. M. A. Levin, X.-G. Wen, String-net condensation: A physical mechanism for topological phases. *Phys. Rev. B* **71**, 045110 (2005).
55. R. Levy, B. K. Clark, Mitigating the sign problem through basis rotations. arXiv:1907.02076 [cond-mat.str-el] (3 July 2019).
56. G. Torlai, J. Carrasquilla, M. T. Fishman, R. G. Melko, M. P. A. Fisher, Wavefunction positivization via automatic differentiation. arXiv:1906.04654 [quant-ph] (11 June 2019).
57. A. J. Kim, P. Werner, R. Valentí, Alleviating the sign problem in quantum Monte Carlo simulations of spin-orbit-coupled multi-orbital Hubbard Models. *Phys. Rev. B* **101**, 045108 (2019).
58. Y. Shoukry, P. Nuzzo, A. L. Sangiovanni-Vincentelli, S. A. Seshia, G. J. Pappas, P. Tabuada, SMC: Satisfiability modulo convex optimization, in *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control, HSCC '17 (ACM, 2017)*, pp. 19–28.
59. S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, B. M. Terhal, The complexity of stoquastic local Hamiltonian problems. *Quant. Inf. Comp.* **8**, 0361–0385 (2008).
60. S. Bravyi, B. Terhal, Complexity of stoquastic frustration-free Hamiltonians. *SIAM J. Comput.* **39**, 1462 (2009).
61. T. Cubitt, A. Montanaro, Complexity classification of local Hamiltonian problems. *SIAM J. Comput.* **45**, 268–316 (2016).
62. T. S. Cubitt, A. Montanaro, S. Piddock, Universal quantum Hamiltonians. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9497–9502 (2018).
63. ISO 16269-4:2010, *Statistical Interpretation of Data—Part 4: Detection and Treatment of Outliers* (International Organization for Standardization, 2010).
64. J. S. Spencer, N. S. Blunt, W. M. C. Foulkes, The sign problem and population dynamics in the full configuration interaction quantum Monte Carlo method. *J. Chem. Phys.* **136**, 054110 (2012).
65. M. J. Wainwright, *High-Dimensional Statistics: A Non-Asymptotic Viewpoint* (Cambridge Univ. Press, ed. 1, 2019).
66. M. Schmidt, G. Fung, R. Rosales, Fast optimization methods for L1 regularization: A comparative study and two new approaches, in *Machine Learning: ECML 2007*, J. N. Kok, J. Koronacki, R. L. de Mantaras, S. Matwin, D. Mladenić, A. Skowron (Lecture Notes in Computer Science, Springer, 2007), pp. 286–297.

Acknowledgments: We are grateful for the many fruitful discussions that have helped shape this work—with A. Werner, M. Schwarz, J. Bermejo-Vega, and C. Krumnow in early stages of the project and, more recently, with M. Troyer, J. Klassen, M. Ioannou, M. L. Baez, H. Pashayan, S. Trebst, A. Kshetrimayum, A. Studt, and A. Nietner. We also thank B. Terhal, M. Ioannou, J. McClean, and M.L. Baez for helpful comments on our draft. **Funding:** D.H., I.R., and J.E. acknowledge financial support from the ERC (TAQ), the Templeton Foundation, the DFG (EI 519/14-1, EI 519/9-1, EI 519/7-1, and CRC 183 in project B01), and the European Union's Horizon 2020 research and innovation programme under grant agreement no. 817482 (PASQuanS). D.N. has received funding from the People Programme (Marie Curie Actions) EU's 7th Framework Programme under REA grant agreement no. 609427. D.N.'s research has been further co-funded by the Slovak Academy of Sciences and by Slovak Research and Development Agency grant QETWORK APVV-14-0878 and VEGA MAXAP 2/0173/17. **Author contributions:** D.H. and I.R. conceived the non-stoquasticity measure and its relation to the

average sign and carried out all analytical and numerical calculations. D.H. and D.N. conceived the complexity proof idea. J.E. contributed to all aspects of this work. All authors participated in discussions and contributed to writing the manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** The python package designed for the numerical simulations is available publicly at (45), making the results shown reproducible. All analytical calculations, in particular the explicit proof of Theorem 2, are presented in the Supplementary Materials.

Submitted 27 March 2020

Accepted 1 July 2020

Published 14 August 2020

10.1126/sciadv.abb8341

Citation: D. Hangleiter, I. Roth, D. Nagaj, J. Eisert, Easing the Monte Carlo sign problem. *Sci. Adv.* **6**, eabb8341 (2020).

Easing the Monte Carlo sign problem

Dominik Hangleiter, Ingo Roth, Daniel Nagaj and Jens Eisert

Sci Adv **6** (33), eabb8341.
DOI: 10.1126/sciadv.abb8341

ARTICLE TOOLS	http://advances.sciencemag.org/content/6/33/eabb8341
SUPPLEMENTARY MATERIALS	http://advances.sciencemag.org/content/suppl/2020/08/11/6.33.eabb8341.DC1
REFERENCES	This article cites 55 articles, 3 of which you can access for free http://advances.sciencemag.org/content/6/33/eabb8341#BIBL
PERMISSIONS	http://www.sciencemag.org/help/reprints-and-permissions

Use of this article is subject to the [Terms of Service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2020 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).