
AFRL-RI-RS-TR-2020-164

# POWER OF NON-STOQUASTIC QUANTUM ANNEALING OPTIMIZATION

INFORMATION SCIENCES INSTITUTE, UNIVERSITY OF SOUTHERN CALIFORNIA

*SEPTEMBER 2020*

FINAL TECHNICAL REPORT

***APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED***

STINFO COPY

**AIR FORCE RESEARCH LABORATORY**
**INFORMATION DIRECTORATE**

■ **AIR FORCE MATERIEL COMMAND** ■ **UNITED STATES AIR FORCE** ■ **ROME, NY 13441**

# NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nations. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RI-RS-TR-2020-164 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /
KRISTI MEZZANO
Work Unit Manager

/ S /
GREGORY HADYNSKI
Assistant Technical Advisor
Computing & Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

# REPORT DOCUMENTATION PAGE

*Form Approved*
**OMB No. 0704-0188**

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
**PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

| 1. REPORT DATE *(DD-MM-YYYY)* | 2. REPORT TYPE | 3. DATES COVERED *(From - To)* |
|---|---|---|
| SEPTEMBER 2020 | FINAL TECHNICAL REPORT | AUG 2018 - MAR 2020 |

**4. TITLE AND SUBTITLE**

Power of Non-Stoquastic Quantum Annealing Optimization

**5a. CONTRACT NUMBER**
N/A

**5b. GRANT NUMBER**
FA8750-18-1-0044

**5c. PROGRAM ELEMENT NUMBER**
62788F

**6. AUTHOR(S)**

Itay Hen

**5d. PROJECT NUMBER**
CYDT

**5e. TASK NUMBER**
US

**5f. WORK UNIT NUMBER**
C2

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
Information Sciences Institute, University of Southern California
4676 Admiralty Way
Marina del Rey CA 90292

**8. PERFORMING ORGANIZATION REPORT NUMBER**

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**

Air Force Research Laboratory/RITQ
525 Brooks Road
Rome NY 13441-4505

**10. SPONSOR/MONITOR'S ACRONYM(S)**
AFRL/RI

**11. SPONSOR/MONITOR'S REPORT NUMBER**
AFRL-RI-RS-TR-2020-164

**12. DISTRIBUTION AVAILABILITY STATEMENT**
Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT**
Despite the excitement brought on by recent technological breakthroughs that have made quantum annealing computing (QuAnCo) optimizers consisting of thousands of quantum bits commercially available, quantum adiabatic protocols have so far failed to deliver on their promise to serve as useful optimizers, i.e., to find bit assignments that minimize the energy, or cost, of discrete combinatorial optimization problems faster than is possible classically. Thus far, no examples, neither experimental nor theoretical, of practical relevance have been found to indicate a superiority of quantum adiabatic optimization (QAO) over traditional methods. The objective of this proposal is to develop a clear understanding of the type of advantages non-stoquastic quantum fluctuations may have to offer in the field of quantum annealing optimization. The effort focuses on comparing the projected performance of quantum adiabatic algorithms equipped with non-stoquastic driving quantum fluctuations with the corresponding performance of their stoquastic, efficiently simulable, analogues and with state-of-the-art conventional solvers running on standard computers.

**15. SUBJECT TERMS**

Quantum annealing; Non-stoquasticity; Quantum adiabatic optimization; Sign problem

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | **KRISTI MEZZANO** |
| U | U | U | UU | 28 | 19b. TELEPHONE NUMBER (*Include area code*) **N/A** |

## Table of Contents

## List of Figures

## 1.0 SUMMARY

Despite the excitement brought on by recent technological breakthroughs that have made quantum annealing computing (QuAnCo) optimizers consisting of thousands of quantum bits commercially available, quantum adiabatic protocols have so far failed to deliver on their promise to serve as useful optimizers, i.e., to find bit assignments that minimize the energy, or cost, of discrete combinatorial optimization problems faster than is possible classically. Thus far, no examples, neither experimental nor theoretical, of practical relevance have been found to indicate a superiority of quantum adiabatic optimization (QAO) over traditional methods.

It is now widely believed that one possibly necessary missing ingredient in the quest for quantum superiority, both within the experimental and the theoretical frameworks of quantum annealing optimization, is the use of so-called `non-stoquastic' Hamiltonians -- Hamiltonians that cannot, or are not known to, be efficiently simulable via classical algorithms such as quantum Monte Carlo (QMC) methods. This is due to the sign problem, which prohibits the efficient simulation of non-stoquastic Hamiltonians by classical means.

The objective of this effort was to develop a clear understanding of the type of advantages that non-stoquastic quantum fluctuations may have to offer in the field of quantum annealing optimization. The effort focused on comparing the projected performance of quantum adiabatic algorithms equipped with non-stoquastic driving quantum fluctuations with the corresponding performance of their stoquastic, efficiently simulable, analogues and with state-of-the-art conventional solvers running on standard computers.

This effort produced four journal articles:

1. Elucidating the Interplay Between Non-stoquasticity and the Sign Problem [1]. The sign problem is a key challenge in computational physics, encapsulating our inability to properly understand many important quantum many-body phenomena in physics, chemistry and the material sciences. Despite its centrality, the circumstances under which the problem arises or can be resolved as well as its interplay with the related notion of `non-stoquasticity' are often not very well understood. We studied the circumstances under which the sign problem emerges and cleared up some of the confusion surrounding this intricate computational phenomenon. We made use of the recently introduced off-diagonal series expansion quantum Monte Carlo scheme with which we analyzed a number of examples that capture the essence of our results.

2. De-Signing Hamiltonians for Quantum Adiabatic Optimization [2]. Quantum fluctuations driven by non-stoquastic Hamiltonians have been conjectured to be an important and perhaps essential missing ingredient for achieving a quantum advantage with adiabatic optimization. We introduced a transformation that maps every non-stoquastic adiabatic path ending in a classical Hamiltonian to a corresponding stoquastic adiabatic path by appropriately adjusting the phase of each matrix entry in the computational basis. We compared the spectral gaps of these adiabatic paths and found numerically that the paths based on non-stoquastic Hamiltonians have generically smaller spectral gaps between the ground and first excited states, suggesting they are less useful than stoquastic Hamiltonians for quantum adiabatic optimization.

3. Quantum Annealing Advantages with Non-stoquastic Hamiltonians [3]. We studied the performance of ideal quantum annealers functioning as estimators of expectation values of physical observables for quantum non-stoquastic Hamiltonians. By contrasting the required effort against that of quantum Monte Carlo algorithms for random Max-Cut instances, we found that an ideal quantum device may provide clear scaling runtime advantages over its classical competitors, which encounter a severe sign problem.

4. Efficient simulations of so-called non-stoquastic superconducting flux circuits [4]. One of the main avenues being explored as of late to obtain quantum speedups by annealing is that of quantum simulations, wherein one performs measurements mid-anneal. To that aim, there is a tremendous ongoing effort to fabricate superconducting flux circuits which are non-stoquastic in their qubit representation, as these circuits are thought to be ``unsimulable'' by classical approaches such as quantum Monte Carlo. We devised a method that successfully simulates in a scalable way the circuit Hamiltonian directly, by utilizing an infinite Hilbert representation that is, in fact, stoquastic.

## 2.0 INTRODUCTION

Despite the excitement brought on by recent technological breakthroughs that have made quantum annealing computing (QuAnCo) optimizers consisting of thousands of quantum bits commercially available [5-9], quantum adiabatic protocols have so far failed to deliver on their promise to serve as useful optimizers, i.e., to find bit assignments that minimize the energy, or cost, of discrete combinatorial optimization problems faster than is possible

classically [10-13]. Thus far, no examples, neither experimental nor theoretical, of practical relevance have been found to indicate a superiority of quantum adiabatic optimization (QAO) over traditional methods.

It is now widely believed that one possibly necessary missing ingredient in the quest for quantum superiority, both within the experimental and the theoretical frameworks of quantum annealing optimization, is the use of so-called `non-stoquastic' Hamiltonians --- Hamiltonians that cannot, or are not known to, be efficiently simulable via classical algorithms such as quantum Monte Carlo (QMC) methods [14-17]. This is due to the sign problem which prohibits the efficient simulation of non-stoquastic Hamiltonians by classical means [18-25].

This effort was devoted to the exploration of the inherent power of QuAnCo systems driven by non-stoquastic quantum fluctuations to accelerate the runtimes to solution of hard optimization problems. So far, attempts to find quantum speedups on problems of this type have focused on using stoquastic driver Hamiltonians to drive the annealing. This is partly because stoquastic fluctuations are easier to simulate on standard computers and also because these type of fluctuations are the only ones that are currently available in experimental settings. To date however, no stoquastically driven speedups have been exhibited.

The present effort focused on comparing the projected performance of quantum adiabatic algorithms equipped with non-stoquastic driving quantum fluctuations with the corresponding performance of their stoquastic, efficiently simulable, analogues and with

state-of-the-art conventional solvers running on standard computers. By studying the minimum gaps of the quantum annealing Hamiltonians we answered the question of whether or not quantum annealing optimization algorithms with non-stoquastic drivers can solve `hard' sets of problems -- those which belong to the NP-complete category and for which all known algorithms take an exponential amount of time (exponential complexity) at least in the worst case.

Previous numerical studies have so far been unable to find indications of quantum speedups on QuAnCo systems. These examined the typical minimum gaps of constraint satisfaction problems and showed that very small gaps could appear in the spectrum of the Hamiltonian due to avoided crossings between the ground state and the first excited state, corresponding to a local minimum of the optimization problem. This `bottleneck' was shown to appear in larger and larger fractions of the instances as the problem size N increases, indicating the existence of a first-order quantum phase transition ultimately leading to an exponentially small typical gap and therefore also to the failure of adiabatic quantum optimization. However, most studies to date have only utilized stoquastic driver Hamiltonians (specifically, the generic transverse-field driver).

To ascertain the prospects of quantum annealing optimization algorithms driven by non-stoquastic drivers, we studied systematically the typical minimum gaps of Hamiltonians interpolating between non-stoquastic drivers and the problem Hamiltonians that encode the optimization problems of interest. We compared the scaling with problem size of the typical minimum gaps of these Hamiltonians against those of the analogous stoquastically driven Hamiltonians. This scaling serves as an indicator of the running times that quantum

annealing optimization algorithms require to solve the input problems with high probability, as promised by the adiabatic theorem of Quantum Mechanics. As we discuss in detail below, we find generally that stoquastic Hamiltonians have larger gaps than non-stoquastic ones. This calls into question the utility of non-stoquastic drivers in quantum adiabatic optimization.

In direct continuation to the above, we continued to examine the ability of quantum annealers to function as `quantum adiabatic simulators' (QASs), i.e., we tested their potential to mimic the operation of quantum Monte Carlo (QMC), that is, to calculate expectation values of observables. QMC is used to simulate systems beyond the reach of exact-diagonalization techniques (under about two dozen particles). However, QMC simulations of non-stoquastic systems are stymied by the infamous sign problem, while a QAS with rich enough interactions is believed not to be afflicted by it. One may therefore ask what sort of advantages, if any, a QAS might offer over standard computers when it comes to simulations of quantum systems and whether it could eventually replace QMC algorithms as the de-facto approach for studying the ground state properties of novel new systems. If indeed QAS devices could offer scaling advantages, QASs will find practical use in a variety of fields from quantum chemistry, material sciences through Condensed Matter Physics to Quantum Chromodynamics and beyond.

We carried out a study to determine the prospects of observing quantum advantages using QASs for sampling from the ground states of quantum non-stoquastic Hamiltonians. For a specific class of non-stoquastic Hamiltonians, we were able to show that QMC runtimes scale exponentially in the inverse-temperature. This to be contrasted with what is expected

from the adiabatic theorem, which dictates that a QAS will only require runtimes scaling as a polynomial in the minimum gap along its interpolating path. This result suggests a clear scaling advantage for QAS over QMC.

## 3.0 METHODS, ASSUMPTIONS AND PROCEDURES

### 3.1 Non-stoquastic quantum adiabatic optimization: Stoquastization of quantum adiabatic Hamiltonians

We introduced a transformation that maps every non-stoquastic adiabatic path ending in a classical Hamiltonian to a corresponding stoquastic adiabatic path by appropriately adjusting the phase of each matrix entry in the computational basis [1]. We compared the spectral gaps of these adiabatic paths and find both theoretically and numerically that the paths based on non-stoquastic Hamiltonians have generically smaller spectral gaps between the ground and first excited states, suggesting they are less useful than stoquastic Hamiltonians for quantum adiabatic optimization. These results apply to any adiabatic algorithm which interpolates to a final Hamiltonian that is diagonal in the computational basis. In this work we define a locality-preserving mapping which takes every non-stoquastic quantum adiabatic optimization protocol to a corresponding stoquastic protocol. Considering various ensembles (dense matrices, signed graphs, and local Hamiltonians) we find that the non-stoquastic adiabatic paths have smaller spectral gaps than the corresponding stoquastic adiabatic paths, with high probability.

### 3.2 Numerical QMC simulations of non-stoquastic Hamiltonians

We carried out a numerical study to determine the prospects of observing quantum advantages using quantum annealing simulators (QASs) for sampling from the ground states of quantum non-stoquastic Hamiltonians -- Hamiltonians with positive off-diagonal elements that cannot, or are not known to, be efficiently simulatable by quantum Monte Carlo algorithms due to the sign problem [3].

For a specific class of non-stoquastic Hamiltonians, we show that QMC runtimes scale exponentially in the inverse-temperature β, and if β must be chosen to be at least proportional to the inverse energy gap of the Hamiltonian in order to sample closely from the ground state, then QMC algorithms will typically scale inversely proportional to the gap. This is to be contrasted with what is expected from the adiabatic theorem, which dictates that a QAS will only require runtimes scaling as a polynomial in the minimum gap along its interpolating path.

To ascertain the computational power of QASs tasked with the estimation of ground-state properties of non-stoquastic Hamiltonians [26,27], we consider a time-dependent Hamiltonian of the form:

$$H(s) = (1-s)H_D + s(1-s)H_E + s\,H_I. \qquad (1)$$

where s is the dimensionless annealing parameter. For H_D, we chose the standard transverse-field $H_D = \sum_i X_i$ [here, $X_i$ ($Z_i$) is the Pauli-x (-z) matrix acting on spin i] whose ground state, the uniform superposition state, is easy to prepare, while for the Ising Hamiltonian $H_I$ we take Max-Cut instances defined on random 3-regular graphs

with Hamiltonian of the form $H_I = \sum_{\langle ij \rangle} Z_i \, Z_j$ (see Fig. 1).

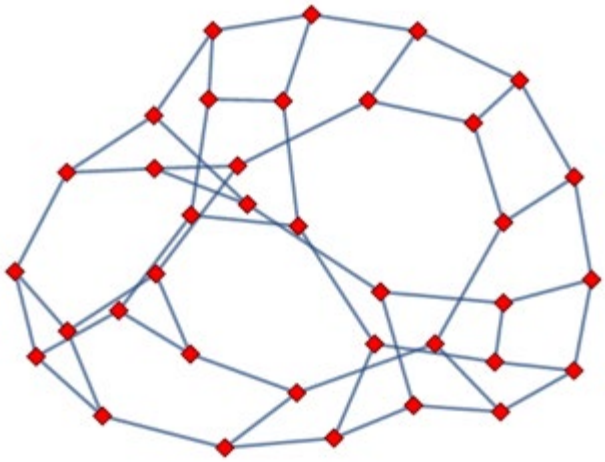

**Figure 1: Connectivity of a randomly generated N=36-spin 3-regular MAX2SAT instance. Here the diamonds denote spins and the edges denote antiferromagnetic couplings with strength J_ij=1. Each spin is connected to three other randomly chosen spins.**

This class of instances corresponds to a particular choice of an Ising Hamiltonian, whereby each spin is coupled antiferromagnetically (with strength $J_{ij} = 1$) with exactly three other spins picked at random and is known to be NP-hard. To incorporate non-stoquasticity we chose the `extra' Hamiltonian $H_E$ to be

$$H_E = \sum_{\langle ij \rangle} X_i \, X_j \tag{2}$$

with the same connectivities ⟨ij⟩ as $H_I$.

The value chosen for the parameter b in the above equation determines whether H(s) is

`standard' (b=0), stoquastic (b=-1) or non-stoquastic (b=1). The b=1 choice is the focus of the present work. The annealing parameter s is varied smoothly from 0 at t=0 to a certain target value $s^* < 1$ at t=T, ensuring that the final Hamiltonian $H(s^*)$ has a non-negligible quantum component (see Fig. 2). This is to be contrasted with standard quantum annealing where the annealing proceeds all the way to s=1 in an attempt to solve an optimization problem encoded in $H_I$. At this point, a computational-basis measurement takes place and the entire process is repeated to gather an appropriate number of samples. Here we study 100 random 3-regular Max-Cut instances at sizes N in the range [6,24]. The target s values we focus on are $s^* = 0.3,0.4,0.5$.

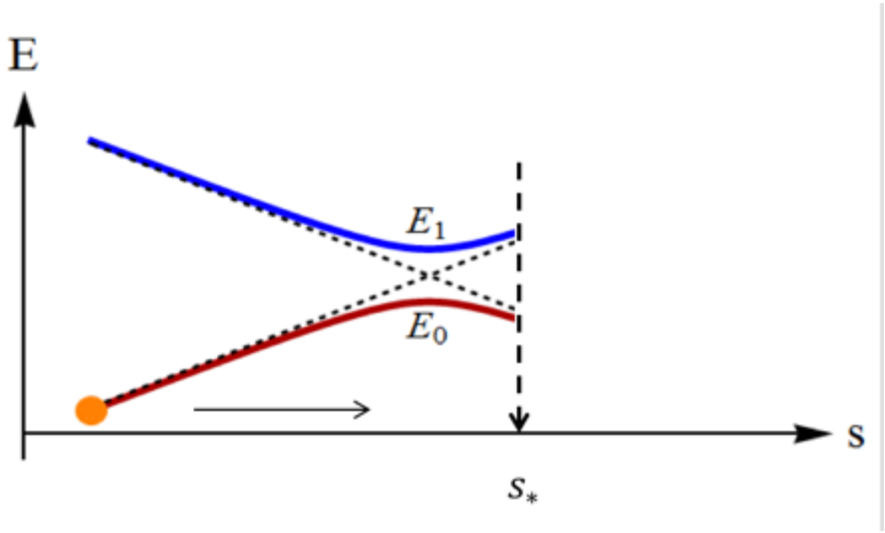


**Figure 2: A quantum adiabatic process with a mid-anneal measurement.**

## 4.0 RESULTS AND DISCUSSION

### 4.1 Gaps of stoquasticized quantum adiabatic Hamiltonians

We demonstrated that the stoquasticized variants of various classes of random matrices become increasingly more favorable to their non-stoquastic versions with increasing system

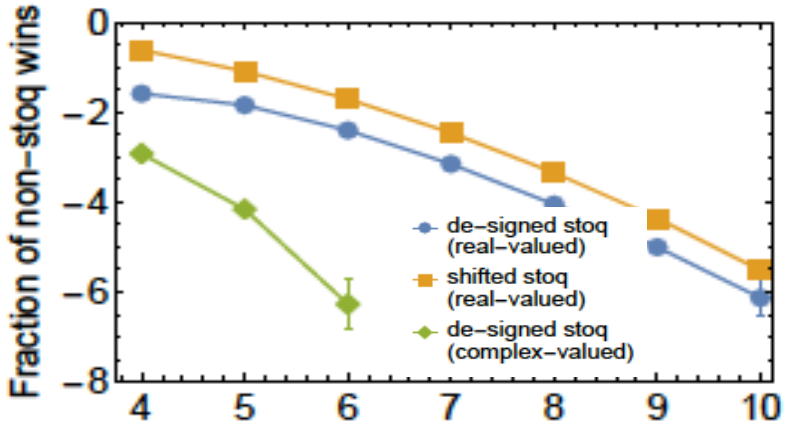


**Figure 3:Fraction of non-stoquastic wins for dense matrices as a function of matrix dimension N. Each data point is the mean over 10 million matrix instances, with the error bars being twice the standard error of the mean.**

size. We studied the gaps of dense and local matrices as well as quantum annealing runtimes for a few classes of problems. In each case we considered both the de-signed and shifted stoquastizations.

We generated random (i) real-valued and (ii) complex-valued Hermitian matrices of different dimensions and compute their gaps. In the real case, matrix entries are drawn independently and uniformly from the range [-1,1] whereas in the complex case the sampling is carried out for the real and imaginary parts separately. To enforce Hermiticity, we average the generated matrices with their conjugate transpositions. As a next step, we

calculated numerically the gaps of the stoquasticized counterparts of the generated matrices. We consider the two forms of stoquastization: de-signed and shifted (see Fig. 3) and Ref.[2].

The figure of merit we focused on is the `fraction of non-stoquastic wins'--- the fraction of occurrences for which the energy gap of the (generally) non-stoquastic matrix is strictly larger than its stoquasticized counterpart. We examined the behavior of this fraction as a function of matrix dimension. The results are summarized in the above figure. As is evident, the fraction decays with the matrix dimension, and already at N=10 the fraction is incredibly small.

We note that the fraction of non-stoquastic wins as measured by the gaps in the figure above does not necessarily translate to the same fraction of wins as measured by an optimized computational cost often referred to as the `average time to solution'. For each instance, we find the annealing time that minimizes this quantity and calculate the median computational cost as a function of system size for the same group of instances. We show our results in Fig. 4 and we see no increase in the fraction of non-stoquastic wins. This discrepancy between the gaps and the computational cost can be attributed to the smallness of the problem instances: these small instances still have relatively large gaps (we still do not see the exponential decay with system size that might be expected for this class of instances), so the gap alone does not quantitatively predict the computational cost.

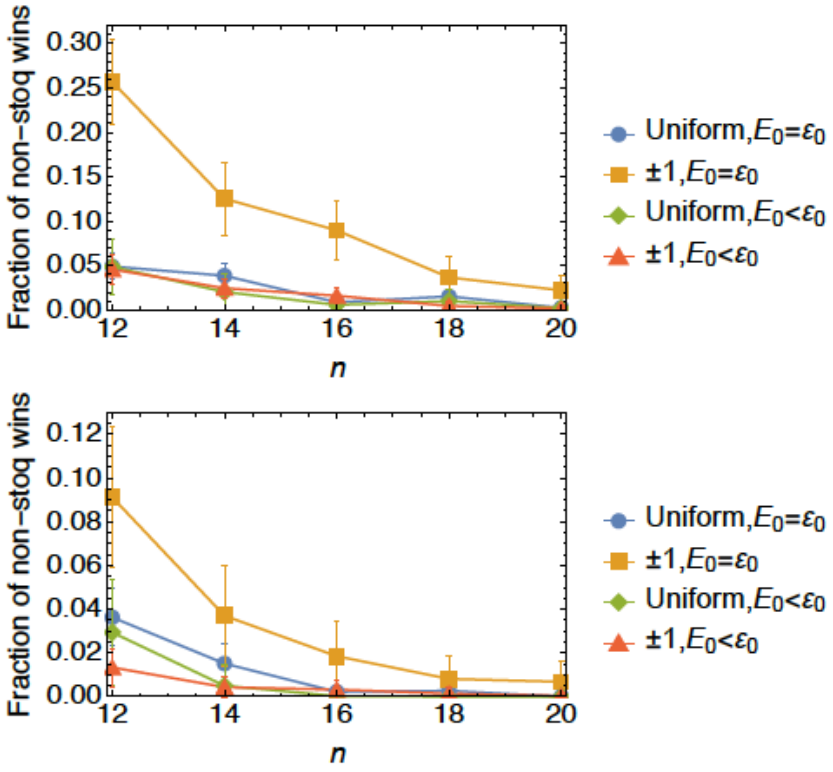


**Figure 4: Fraction of non-stoquastic occurrences of smaller non-stoquastic average time to solution relative to their stoquasticized analogues as a function of problem size, for Max-Cut problems with an interpolating Hamiltonian. Error bars correspond to two standard deviations confidence calculated with a bootstrap over 1000 instances**

### 4.2 Quantum annealers as quantum simulators

We studied the behavior of a typical small instances (N=6...14) as a function of the inverse-temperature β of the QMC algorithm [28-29], to compare against the projected performance of a quantum annealer simulator (QAS).

The partition function associated with the non-stoquastic quantum Hamiltonians cannot (or is not known to) be broken up into positive-valued Monte Carlo weights. The common

practice to tackle this is by using the absolute value of the weights to determine the probabilities associated with the Markov process. For models that do not have a sign problem, all weights are positive. For non-stoquastic systems, negative weights are equally dominant and, in this case, thermal averages of physical observables will fluctuate rapidly, resulting in extremely large error bars and requiring an exponentially large number of measurements. An appropriate figure of merit for how adverse the sign problem is for a given QMC algorithm is given by the `weighted sign' $\langle sgn \rangle$: the ratio of the sum of all weights divided by the sum of their absolute values. Its value is inversely proportional to the QMC runtime (see Ref.[1] for more details).

We studied the behavior of a typical small instances (N=6…14) as a function of the inverse-temperature β of the QMC algorithm. For our QMC algorithm, we use the ``off-diagonal expansion" QMC algorithm [1], which expands the partition function in terms of products of the off-diagonal operators. We find that even these small sizes pose severe challenges to QMC which does not allow studying larger problem sizes or colder temperature

As shown in Figure 5, we find an exponential growth in the number of required samples with increasing inverse temperature β and with problem size N. This is associated with an exponential decrease in the weighted sign.

We note that even these small sizes pose severe challenges to QMC which does not allow studying larger problem sizes or colder temperature. As shown in Fig. 5 we find an exponential growth in the number of required samples with increasing inverse temperature β and with problem size N at a fixed s*.

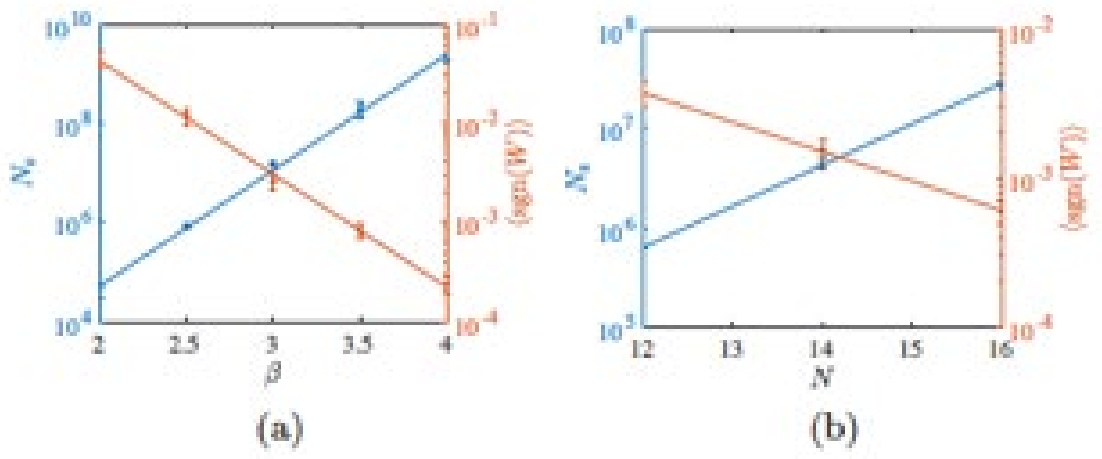


**Figure 5: QMC sign-problem scaling. (a) Number of QMC samples and typical expectation value of the weighted sign as measured with the number of samples required to achieve p=0.1 for a typical N=6 instance as a function of inverse temperature. (b): Number of samples required to achieve p=0.1 and weighted sign as a function of problem size (here, β=2 =. Solid lines are guides for the eye. Error bars correspond to two standard deviations statistical confidence.**

We contrasted the typical runtime scaling with problem size of QMC and QAS, in Figure 6. For comparison, we also include the runtime scaling of the Lanczos exact diagonalization algorithm. Exact diagonalization scales exponentially with problem size n but only as a power-law in the inverse gap. Unlike QMC however, the memory requirements of the Lanczos algorithm are exponential in n whereas those of QMC are linear in n. As is evident from the figure, the QAS runtimes scale far more mildly than both QMC and exact-diagonalization, assuming the ground state can be reached by QAS. These results are positive indications of a quantum scaling advantage over these algorithms for the study of non-stoquastic quantum systems.

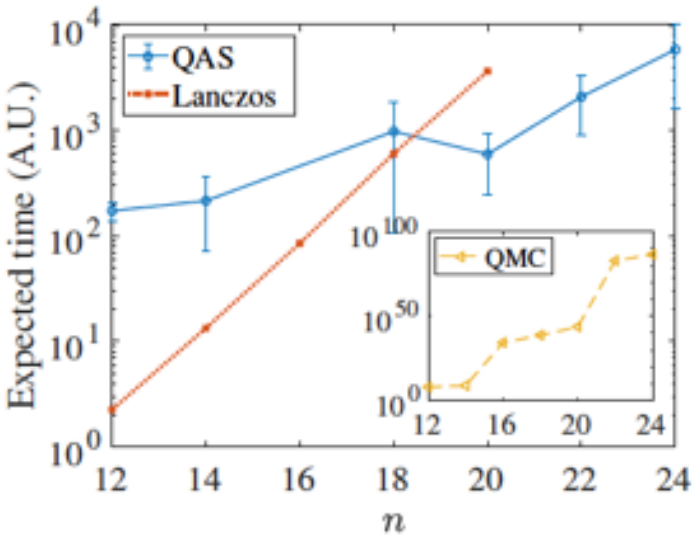


**Figure 6: Typical runtime scaling of QAS, QMC and Lanczos (exact diagonalization), median runtime over all instances as a function of system size n. The error bars correspond to one standard deviation confidence intervals calculated over the instances.**

## 5.0 CONCLUSIONS

In this effort we made an attempt to elucidate the QMC sign problem, one of the most fundamental bottlenecks of many-body physics simulations and alleviate the confusion surrounding certain aspects of the problem. Along the way, we derived a formalism that allowed us to characterize the conditions under which the sign problem emerges and to analyze several core examples that illustrate a number of important observations. First and foremost, we demonstrated that non-stoquasticity --- the non-positivity or complex-valuedness of the off-diagonal entries of the matrix representation of Hamiltonians --- does not imply the existence of a sign problem. We have shown rather that the emergence of a sign problem is tightly connected to the cumulative phase of products of off-diagonal terms along closed paths in the hypercube of basis states. We have also shown that a single spin-

1/2 particle cannot possess a sign problem but that a single spin-one particle can. We have additionally provided an example illustrating that a physical model may be non-stoquastic and have a severe sign problem despite having all-positive ground-state amplitudes, a property that is usually linked with stoquasticity.

Developing a true understanding of the nature of the QMC sign problem is important in physics, chemistry, the material sciences and well beyond those, and is crucial to the potential resolution of the problem. We therefore hope that our work will provide a useful framework for studying the nature of the sign problem in models of physical relevance beyond those analyzed here and will allow addressing the sign problem in more general settings. Another area in which developing a true understanding of the problem is important is the field of quantum computing. This is because quantum simulations or simulations of quantum many-body systems on quantum computers, as originally envisioned by Richard Feynman is one of the most promising applications of near-term quantum devices (as well as the more distant fault-tolerant universal quantum computers). The existence of quantum simulation speedups hinges on the premise that simulating quantum systems is an intractable task for standard computers and relies (at least in part) on the intractability of the sign problem. In this context, it is worth noting in particular the common misconception that the resolution of the sign problem would imply in general a polynomial-time equilibration of QMC simulations or that it would somehow provide as a result a resolution to the famous P versus NP problem of computer science.

We also provided analytical as well as numerical evidence in favor of the assertion that non-stoquastoic driver Hamiltonians are inferior to their stoquastic variants for quantum

annealing optimization tasks. We analyzed the gaps of several types of random non-stoquastic Hamiltonians comparing them against their `stoquasticized' counterparts. We find that generically the non-positivity of the latter Hamiltonians renders their gap larger than that of non-stoquastic ones, making non-stoquastic Hamiltonian less favorable for quantum annealing optimization. Our results imply that stoquastic Hamiltonians should generally be preferable to non-stoquastic ones, at least as far as runtimes of quantum adiabatic algorithms are concerned. It should be noted that examples to the contrary nonetheless exist. While the above examples provide flashes of optimism about the prospects of non-stoquastic drivers in QAO, our results call into question the promise attributed to non-stoquastic drivers to serve as generic catalysts of quantum speedups.

We have compared the effort required from quantum adiabatic simulators and its classical competitor, the quantum Monte Carlo algorithm, to estimate ground-state properties of non-stoquastic quantum Hamiltonians. We have shown that the number of samples, and hence the runtime, required by QMC scales exponentially with system size and inverse gap, while QAS runtimes scale only polynomially with the inverse minimum gap. We have thus demonstrated that QASs may very well function as promising candidates to study non-simulable quantum systems in that they offer a practical and more efficient way to avoid the sign problem.

While promising, our work does not perform a complete survey of classical algorithms that can be used to study the ground state of non-stoquastic Hamiltonians. Depending on the accuracy required in the simulation, tensor network based methods can be used. These methods are unaffected by the sign problem and are polynomially efficient in problem size

but scale exponentially with the `bond dimension'. We leave the question of how the bond dimension needs to scale with problem size to maintain a suitable level of accuracy. It would be interesting to see how bond dimension must scale with problem size to maintain the required accuracy to accurately reproduce ground state expectation values of non-stoquastic Hamiltonians. Another candidate classical algorithm is to use an artificial neural network as an ansatz for the ground state wavefunction. How the number of nodes in the neural network and the number of steps needed to 'train' it to achieve the necessary accuracy for the ground state of non-stoquastic Hamiltonians is an important question that we leave for future study.

## 6.0 REFERENCES


1. L. Gupta and I. Hen, Advanced Quantum Technologies. arXiv:1910.13867 (2019).
2. E. Crosson, T. Albash, I. Hen, and A. P. Young, submitted for publication. arXiv:2004.07681 (2020).
3. T. Albash, A. P. Young and I. Hen, Quantum Annealing Advantages with Non-stoquastic Hamiltonians, in preparation.
4. T. Halverson, L. Gupta, M. Goldstein and I. Hen, Efficient simulations of so-called non-stoquastic superconducting flux circuits, in preparation.
5. M. W. Johnson, M. H. S. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. S. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, Nature 473, 194 (2011).
6. J. Brooke, D. Bitko, T. F., Rosenbaum, and G. Aeppli, Science 284, 779 (1999).
7. T. Kadowaki and H. Nishimori, Phys. Rev. E 58, 5355 (1998).
8. E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, Science 292, 472 (2001).
9. G. E. Santoro, R. Marto_n_ak, E. Tosatti, and R. Car, Science 295, 2427 (2002).
10. A. P. Young, S. Knysh, and V. N. Smelyanskiy, Phys. Rev. Lett. 101, 170503 (2008).

11. A. P. Young, S. Knysh, and V. N. Smelyanskiy, Phys. Rev. Lett. 104, 020502 (2010).

12. I. Hen and A. P. Young, Phys. Rev. E 84, 061152 (2011).

13. E. Farhi, D. Gosset, I. Hen, A. W. Sandvik, P. Shor, A. P. Young, and F. Zamponi, Phys. Rev. A 86, 052334 (2012), (arXiv:1208.3757).

14. T. F. Ronnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, Science 345, 420 (2014).

15. I. Hen, J. Job, T. Albash, T. F. R_nnow, M. Troyer, and D. A. Lidar, Phys. Rev. A 92, 042325 (2015).

16. S. Boixo, T. Albash, F. M. Spedalieri, N. Chancellor, and D. A. Lidar, Nat. Commun. 4, 2067 (2013).

17. T. Albash, W. Vinci, A. Mishra, P. A. Warburton, and D. A. Lidar, Phys. Rev. A 91, 042314 (2015).

18. M. Troyer and U.-J. Wiese, Phys. Rev. Lett. 94, 170201 (2005).

19. E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Phys. Rev. B 41, 9301 (1990).

20. H. Nishimori and K. Takada, Frontiers in ICT 4, 2 (2017).

21. L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, Phys. Rev. B 95, 184416 (2017).

22. E. Crosson, E. Farhi, C. Yen-Yu Lin, H.-H. Lin, and P. Shor, ArXiv e-prints (2014),

arXiv:1401.7320 [quant-ph].

23. M. K. Garey and D. S. Johnson, Computers and Intractability. A Guide to the Theory of NP-Completeness (H. H. Freeman, New York, 1997).

24. S. Bravyi and B. Terhal, SIAM Journal on Computing, SIAM Journal on Computing 39, 1462 (2009).

25. M. Marvian, D. A. Lidar, and I. Hen, ArXiv e-prints (2018), arXiv:1802.03408 [quant-ph].

26. T. Kato, J. Phys. Soc. Jap. 5, 435 (1950).

27. S. Jansen, M.-B. Ruskai, and R. Seiler, J. Math. Phys. 48, (2007).

28. T. Albash, G. Wagenbreth, and I. Hen, Phys. Rev. E 96, 063309 (2017).

29. C.-W. Liu, A. Polkovnikov, and A. W. Sandvik, Phys. Rev. Lett. 114, 147203 (2015).

## List of Symbols, Abbreviations and Acronyms

QAS - Quantum Annealing simulator

QAO - Quantum Annealing Optimization

QuAnCo - Quantum annealing computing

QMC - Quantum Monte Carlo