Distribution Agreement

In presenting this thesis as a partial fulfillment of the requirements for a degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis in whole or in part in all forms of media, now or hereafter now, including display on the World Wide Web. I understand that I may select some access restrictions as part of the online submission of this thesis. I retain all ownership rights to the copyright of the thesis. I also retain the right to use in future works (such as articles or books) all or part of this thesis.

Justin Y. J. Burton April 13, 2021

A New Method for Heuristic Evaluation by Means of Finite Size Scaling

by

Justin Y. J. Burton

Stefan Boettcher
Adviser

Physics

Stefan Boettcher

Adviser

Justin C. Burton

Committee Member

Michelangelo Grigni

Committee Member

Effrosyni Seitaridou

Committee Member

2021

A New Method for Heuristic Evaluation by Means of Finite Size Scaling

By

Justin Y. J. Burton

Stefan Boettcher

Adviser

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Physics

2021

Abstract

A New Method for Heuristic Evaluation by Means of Finite Size Scaling
By Justin Y. J. Burton

We develop a new method for evaluating the efficacy of NP-hard optimization heuristics as an alternative to testbed heuristic evaluation. The new method evaluates heuristics using finite-size scaling by extrapolating the observed relationship between resource utilization and maximum system size with optimized solutions to larger system sizes. This improves upon testbed heuristic evaluation by providing better evaluation of average-case problems and providing insight into proper resource allocation to efficiently find heuristic solutions. We demonstrate this on the parallel tempering algorithm for the Edwards-Anderson model of a hypercubic spin-glass with periodic boundary conditions and bimodal bond distribution. We run the parallel tempering algorithm with different time resources and observe the system sizes where the algorithm deviates from the accepted ground state energies. We observe deviations between systems sizes with 343 spins and 729 spins for between 125 and 3000 time steps of the parallel tempering algorithm. This information is extrapolated to evaluate the efficacy of the parallel tempering algorithm at finding ground states of larger spin-glasses.

A New Method for Heuristic Evaluation by Means of Finite Size Scaling

By

Justin Y. J. Burton

Stefan Boettcher

Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Science with Honors

Physics

2021

Table of Contents

List of Figures

List of Tables

A New Method for Heuristic Evaluation by Means of Finite Size Scaling

Justin Burton

Stefan Boettcher, Ph. D.

Emory Department of Physics

April 2021

1 Introduction

1.1 NP-Hard Problems and Heuristics

Important problems in physics and computer science including the Boolean satisfiability problem, the traveling salesman problem, the spin glass ground state problem, and neural network training belong to the NP-hard class of problems [2][3]. If $P \neq NP$, the NP-hard problems do not have polynomial-time algorithms to find exact solutions. In the absence of a polynomial-time algorithm, heuristics become necessary to solve large NP-hard problems. A polynomial-time algorithm is an algorithm with a polynomial function of parameter size describing execution time. This is in contrast to an exponential-time algorithm, which has an execution time which scales exponentially with parameter size. Heuristic solutions must come with limitations compared to an ideal polynomial-time exact solution, and different heuristics for the same problem can posses different limitations. When comparing two heuristics, one may offer a superior approximate solution to the other, but at the cost of increased resource

utilization. Given this, evaluating the efficacy of heuristics should provide insight into the advantages and disadvantages provided by a heuristic compared to its alternatives.

1.2 Heuristic Evaluation Methods

Testbeds are the dominant method for evaluating the efficacy of optimization heuristics. Testbeds provide a heuristic with sample problems often over representing worst-case problems. While effective for demonstrating the worst-case efficacy of a heuristic, this can be unhelpful when presented with the more likely average-case behavior, upon which heuristics can perform far better. Furthermore, the same testbeds being used over decades results in heuristic development prioritizing testbeds evaluation over other virtues of a heuristic. This hurts the quality of testbed evaluation as a measure. Additionally, testbed evaluation may not provide enough information to understand under what conditions certain heuristics outperform others. With more information, one could know how much resources a heuristic demands when solving a specific problem [4].

In this work we develop a new method of heuristic evaluation utilizing the finite size scaling method of extrapolation used in statistical mechanics. Finite size scaling observes the behavior of optimized solutions to statistical systems at small system sizes then extrapolates this behavior to the thermodynamic limit. Using our new method, we first find the relationship between resource parameters and maximum system size with optimized solutions. After observing the relationship between resources and solution optimization, we extrapolate this behavior using finite size scaling to form an evaluation of the heuristic. This method can be repeated for different resource parameters of the heuristic. We may also adjust qualitative parameters in heuristics and observe how they change the relationship between resource utilization and solution optimization.

This method of heuristic analysis was previously utilized to measure the performance of iterated tabu search on the Quadratic Unconstrained Binary Optimization (QUBO) problem [5]. This work seeks to further investigate and standardize the method of heuristic analysis

Figure 1: These are diagrams of a two-spin Ising model. Filled circles represent a spin of $S_i = +1$ while open circles represent a spin of $S_i = -1$. If we assume the interactions between our spins is $J_{ij} = 1$, a ferromagnetic interaction, the system on the left has total energy $H = -1$ while the system on the right has total energy $H = 1$ based on Equation 1. In this case, the diagram on the left represents a ground state of this system. For larger Ising models, it remains true that two like spins with ferromagnetic bond interaction $J_{ij} > 0$ will lower the total energy of the system, while opposite spins will increase the total energy. The opposite is true for anti-ferromagnetic bond interactions $J_{ij} < 0$.

by testing the method on a new model and algorithm.

1.3 The Edwards-Anderson Spin Glass

In this work, we investigate a heuristic built to find the ground state of the Edwards-Anderson spin glass, a variation of the Ising spin glass. A spin glass is a magnetic state characterized by random alignment of interactions between particle spins. The bond interaction J_{ij} between two particles i and j defines the energy of the system that contains the two spins. For example, in a ferromagnet $J_{ij} > 0$ for all non-zero interactions between i and j, leading to a ground state of all spins pointing in the same direction. This bond is a ferromagnetic bond, and the opposite $J_{ij} < 0$ represents an antiferromagnetic bond.

A spin glass is expected to have an equal number of ferromagnetic and antiferromagnetic interactions, resulting in insignificant total magnetization even at low temperature. Another feature of spin-glasses is a rough energy landscape as a result of frustrated plaquettes: regions in which a spin locally does not favor either +1 or -1. An example of a frustrated plaquette is illustrated in Figure 4. While locally the value of a spin may not change its total energy, its spin may affect other frustrated plaquettes resulting in a very different total energy. A result of this is that a spin glass is never observed to be in thermal equilibrium, indicating that algorithms mimicking real thermal easing may fail to reach the ground state configuration [8].

While there exists polynomial-time algorithms for the two-dimensional Ising spin glass,

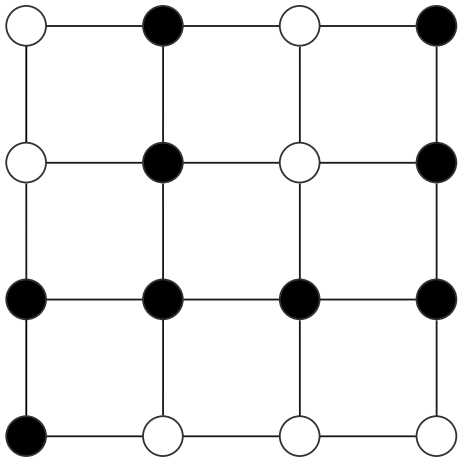

Figure 2: Combining spins and bonds as shown in Figure 1 into a square lattice, we have a two-dimensional Ising model. The lines connecting circles in this diagram represent non-zero bond interactions. The Edwards-Anderson spin glass we investigate in this project only contains non-zero interactions when spins are neighbors as shown here. The spins in this diagram are represented by filled and open circles as described in Figure 1. If the bond interactions are all ferromagnetic, the ground state of this system would have $S_i = S_j$ for all i and j.

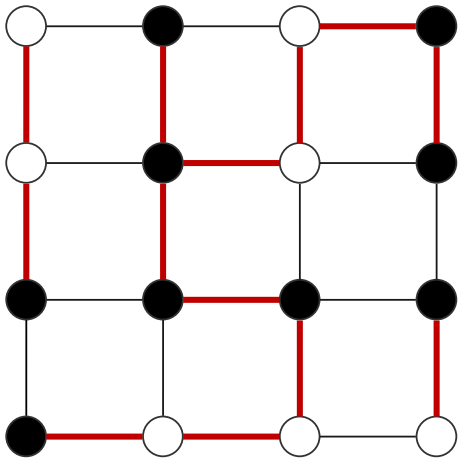

Figure 3: This diagram expands upon the two-dimension Ising model in Figure 2 by specifying bimodal bond distribution. Bond interactions are assigned $J_{ij} = +1$, represented by a red thicker line, or $J_{ij} = -1$, represented by a black thinner line, each with probability 0.5. This random distribution of ferromagnetic bonds and anti-ferromagnetic bonds results in a spin-glass, the system type we optimize in this work. Because only neighboring bonds have non-zero interaction, this is a diagram of a two-dimensional Edwards-Anderson model with bimodal bond distribution.

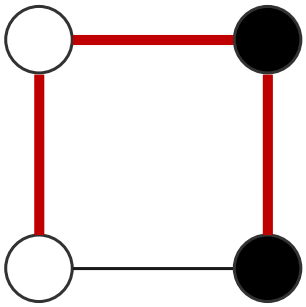

Figure 4: This is the four spins in the upper-right corner of the Edwards-Anderson spin-glass shown in Figure 3. A section of spins such as this is called a plaquette. The upper-left spin of this plaquette can change its value without affecting the total energy of the plaquette. The change in total energy contributed by one bond is reversed by the energy change contributed by the other bond. Notice, however, that the upper-right spin shares this property while the lower-left spin does not. If we swap the spin of the upper-left spin, the lower-left spin now has no effect on total energy, while the upper-right spin does. These properties indicate a frustrated plaquette. While the energy of this plaquette is unaffected by certain spin swaps, the effects on the total energy of the system can be large. This is a cause of the rough energy landscape of spin glasses.

the three-dimensional Ising spin glass is NP-hard. This is due to the existence of a reduction of finding the cocycle of maximum cardinality in a cubic graph to the simplest three-dimensional Ising spin-glass problem: Finding the ground state of a two-level Ising spin-glass with bimodal bond interactions [2].

Our model is a hypercubic lattice of spins with arbitrary dimension and side-length, however due to the NP-hard nature of spin-glasses with dimension greater than two, we will mostly study the three-dimensional case. The Edwards-Anderson model only considers the interactions of neighboring spins and calculates total energy according to Equation 1 and magnetization according to Equation 2, where H is the energy of the system, m is the magnetization of the system, n is the number of spins in the system, J_{ij} is the interaction between spins i and j, and S_i is the spin at index i. Our model uses bimodal bond distribution, a spin-glass interaction in which each non-zero interaction is chosen to be +1 or -1 with probability 0.5 each. The result is a roughly equal number of ferromagnetic and antiferromagnetic interactions. Our model uses periodic boundary conditions to mitigate the relatively greater effects fixed boundary conditions would have on smaller systems compared

to the more realistic very large systems we hope to investigate [9].

$$H = -\sum_{\langle ij \rangle} J_{ij} S_i S_j \tag{1}$$

$$m = \frac{1}{N} \sum_i S_i \tag{2}$$

While the study of spin-glasses is an active and important problem in the area of Statistical Mechanics, its heuristics are important to fields such as operation research and financial asset management due to the spin-glass problem being equivalent to the quadratic unconstrained binary optimization (QUBO) problem. This equivalency opens the heuristics for the spin-glass problem to solve a wide array of problems with applications in many fields [5]. Additionally, the spin glass provides the simplest model of a glassy system, a category of systems with similar behaviors as a result of disordered constructions [11].

1.4 The Parallel Tempering Algorithm

The heuristic we are using is the Parallel Tempering algorithm, an improvement on the Metropolis-Hastings algorithm, a Markov chain Monte-Carlo algorithm [10]. Our applications of the Metropolis-Hastings algorithm begins with a random configuration of spins. To generate a new configuration, a random spin is chosen and the configuring with this spin flipped is investigated. If the new configuration has lower energy than the last, the flip is accepted immediately into the new configuration. If the new configuration has higher energy than the last, the flip is accepted with a probability given by the Boltzmann distribution as shown in Equation 3. This process is repeated n times to produce one sweep of the Metropolis algorithm [9].

$$P(s \to s') = \begin{cases} 1, & \text{if } \Delta H \leq 0 \\ e^{\frac{-\Delta H(s,s')}{kT}}, & \text{if } \Delta H \geq 0 \end{cases} \tag{3}$$

The parallel tempering algorithm improves upon the Metropolis-Hastings algorithm by providing a method for models to escape local minima. The parallel tempering algorithm escapes local minima by running many metropolis algorithms in parallel at a distribution of temperatures, then choosing to swap the temperatures between lattices of neighboring temperatures with the probability given in Equation 4 based on the Boltzmann distribution. The resulting variation in temperature gives spin configurations in local minima an opportunity to escape that minimum during the next sweep of the metropolis algorithm [7].

$$P[(E_i, T_i) \to (E_{i+1}, T_{i+1})] = \min[1, \exp[(E_{i+1} - E_i)(1/T_{i+1} - 1/T_i)]] \tag{4}$$

2 Methods

2.1 Language and Methods for Algorithm Development

The model, heuristic, and evaluation were coded in the C programming language compiled using the GNU Compiler Collection (GCC). The GNU Scientific Library's Taus random number generator was used for random number generation, which is necessary for the stochastic processes involved both in the construction of our models and the heuristic used to find optimized solutions. The random number generators used in the generation of data were seeded with the C library function `time(0)`.

The information for our model is stored on a one-dimensional array, regardless of dimension. This is made possible by an algorithm developed to find the neighbors of a multi-dimensional lattice when spins are indexed linearly. Every element of the one-dimensional array stores a list of neighbors and a list of bond interactions. The bond interactions were given values of -1 or +1 with probability 0.5 for each, as required by bimodal bond distribution. This spin configurations for this model are stored on another one-dimensional array with each element holding the value of the spin. The spins were also given values of -1 or +1 in a similar manner to the bond interactions. While more space efficient methods of

storing an array of bits exist, the memory requirements of the parallel tempering heuristic are inconsequential compared to the time resources, and thus we preferred the time saved by storing computation-ready integers.

The parallel tempering algorithm requires the generation of a large number of spin configurations. We calculate the Hamiltonian of each configuration according to Equation 1 storing the total energy of each configuration in an array. Storing the total energies improves the speed of the algorithm, as the local energy changes from the Metropolis-Hastings algorithm can be easy calculated and used to adjust the new total energy.

2.2 Distribution of Temperatures

Acceptance probabilities approximately independent of temperature are desirable to avoid temperature-space segmentation, and for this reason a non-linear temperature distribution was chosen which has greater density of temperatures near the minimum temperature of the distribution [9]. The temperatures were distributed based on a minimum temperature $T_1 = 0.1$ and maximum temperature $T_M = 2.1$. This was found to have a sufficient number of temperatures above and below the known transition temperature between a spin glass and thermal disorder, $T_C \approx 1.102$ [1]. The resulting distribution is defined by the sequence defined in Equation 5. Our algorithm utilizes 50 parallel spin configurations and thus we generate a distribution of 50 temperatures and we attempt a parallel tempering swap 50 times before returning to the parallel metropolis algorithms. Parallel tempering swaps are implemented simply by swapping the pointers in the array of spin configurations. Each element in our array is associated with a temperature in our distribution, so the effect is that the temperatures associated with the two configurations are swapped.

$$T_k = T_1 \prod_{i=1}^{k-1} R_i \qquad\qquad R_i = \sqrt[M-1]{\frac{T_M}{T_1}} \tag{5}$$

2.3 Multithreaded Programming

One parallel tempering time step consists of one metropolis algorithm sweep for each spin configuration and 50 randomly selected parallel tempering swap attempts. A convenient result of running the metropolis algorithm in parallel is the ability to use multi-threading. Because the process of the Metropolis-Hastings algorithm is identical for each spin configuration with the sole exception of temperature, and no memory is shared between the individual Metropolis algorithms, this process is embarrassingly parallel. A separate POSIX thread was created for each spin configuration, and once the Metropolis algorithm sweep was complete for each configuration the threads were joined together for parallel tempering swaps.

2.4 Finding Ground States

After a set number of parallel tempering time steps, the lowest energy configuring found is the approximate solution found in our heuristic. While there is variation in the lowest energy of the ground state due to the randomly generated spin-glass models, the mean value of repeated tests is predictable. In our three-dimensional cubic lattice, we expect the mean ground state energy to follow $e_3(n) = e_3(\infty) + A/n$ for n-spin lattices [6]. Given the expected results, we repeat the above steps with different time resources in the form of variable parallel tempering steps to see at what system size n our heuristic averages energy that is significantly higher than $e_3(n)$, indicating the heuristic frequently failed to reach the ground state.

A relationship can then be found between the system size at which the heuristic fails to optimize the spin configurations and the time resources allocated. This relationship can be extrapolated to larger system sizes to give a prediction of heuristic behavior in the average case depending on time resources.

3 Results

We plot the average energy per spin $\langle E_n \rangle = E_{total}/n$ vs the reciprocal of the system size $1/n$. The results until a point of upward inflection match the known ground state average energies well [6]. The points of upward inflection appear at larger system sizes as the number of parallel tempering time steps increase.

The plots in Figure 4 show the relationship between $\langle E_n \rangle$ and $1/n$ for the $\langle E_n \rangle$ for ground states proposed by the parallel tempering algorithm with differing time resources measured by number of parallel tempering sweeps. We define the smallest n at which $\langle E_n \rangle < \langle E_{n+1} \rangle$ to be the maximum system size of the heuristic with the given time resources with feasible results, $n_{max}(t)$. While this point may deviate from the known mean energies of the ground

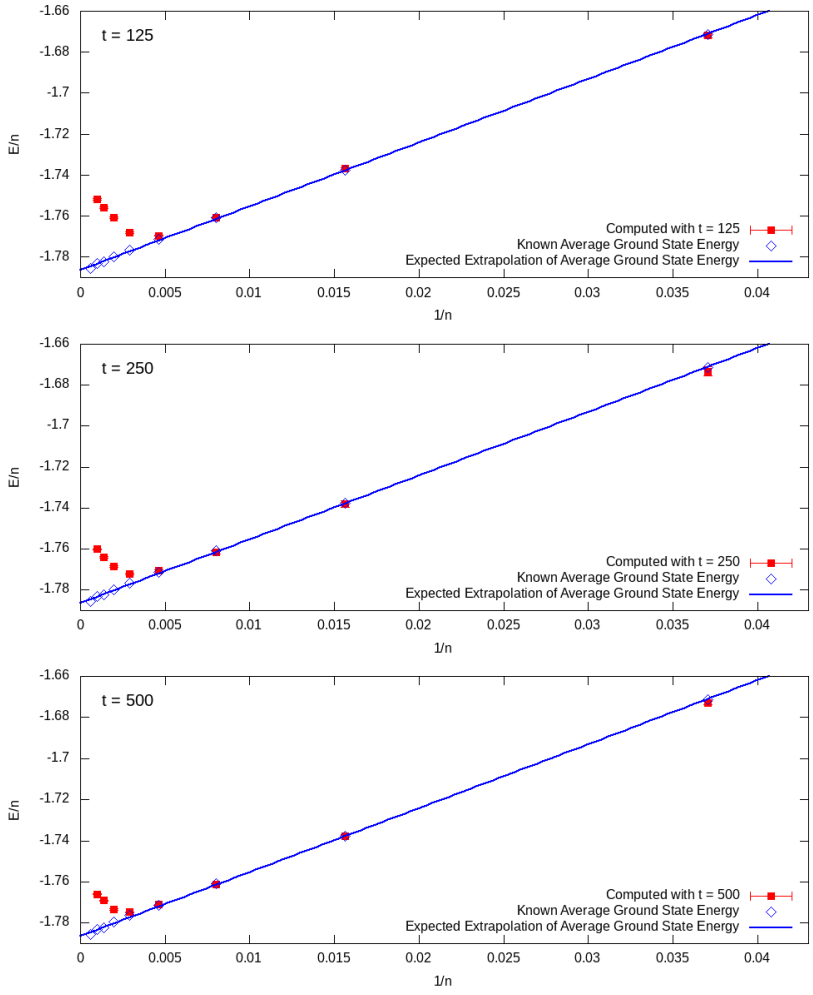

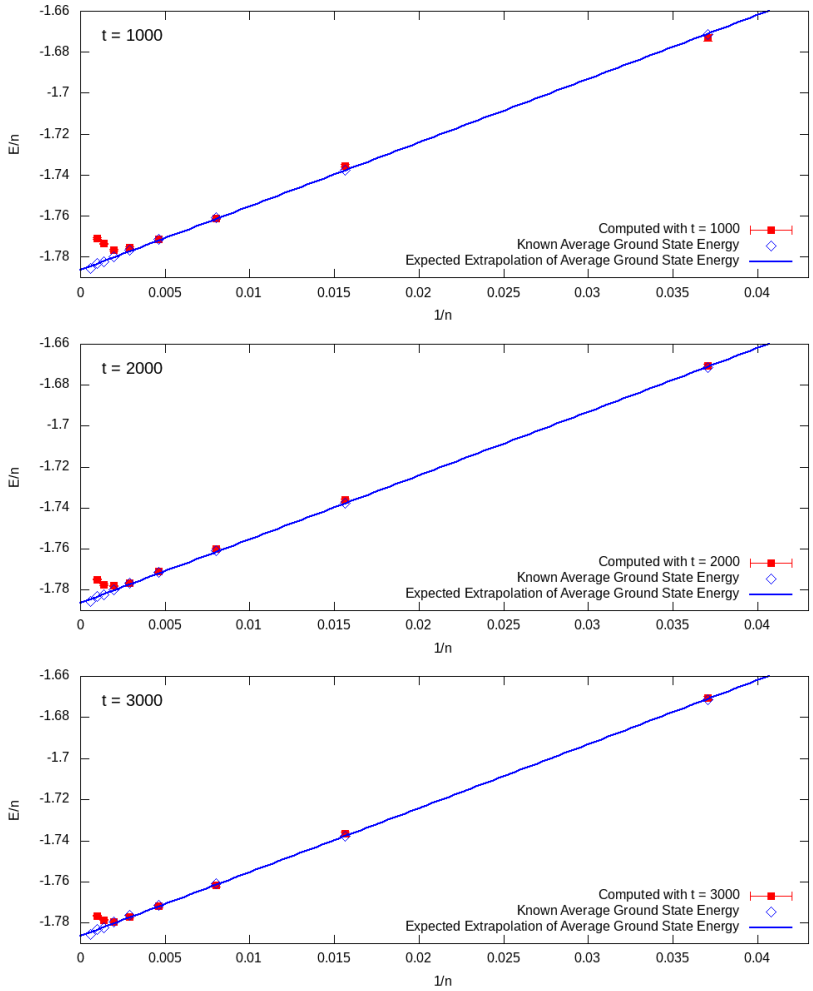

Figure 4: This figure showcases 6 plots comparing the mean Energy per spin $\langle E_n \rangle = E_{total}/n$ of the proposed ground state found by the parallel tempering algorithm with the reciprocal of the system size $1/n$. In the order of the plots from the top to the bottom, the time resources available are $t = 125$, $t = 250$, $t = 500$, $t = 1000$, $t = 2000$, $t = 3000$ in units of parallel tempering time steps. The blue line and points represent the known mean energy per spin of the ground states of the Edwards-Anderson Spin Glass and a linear fit [6]. Each plot showcases a point of deviation at which the points found by our heuristic deviate from accepted values significantly. We denote the smallest n at which $\langle E_n \rangle < \langle E_{n+1} \rangle$ as the maximum system size of the heuristic with the given time resources with feasible results.

state, we decide the point at which $\langle E_n \rangle > \langle E_{n-1} \rangle$ is a point at which the heuristic could no longer be producing fully optimized results.

We find rate of growth of $n_{max}(t)$ with respect to t to decrease with time, as $n_{max}(1000) = n_{max}(2000) = n_{max}(3000)$. This demonstrates a large range where $n_{max} = 9$ compared to the considerably smaller range where $n_{max} = 8$, which is at most $\Delta t = 875$.

Table 1: n_{max} for each tested t

t	$n_{max}(t)$
125	343
250	512
500	512
1000	729
2000	729
3000	729

4 Conclusions

We find a positive correlation of $n_{max}(t)$ with respect to t. We find that the necessary increase in t to increase $n_{max}(t)$ appears to increase with t.

We hope to interpolate our results for E_n with respect to t to find the minimum and utilize finite-size scaling to make predictions about where this minimum would fall given greater time resources. Given the known results that E_n should decrease with n, the n which gives the minimum E_n should indicate the maximum n for which t parallel tempering time steps could be sufficient time for our heuristic to optimize the system. Proper finite-size scaling analysis will likely require more time at different system sizes.

We could further analyze the parallel tempering heuristic by observing and extrapolating the effect of the number of temperatures in our distribution on the ability or time needed for the heuristic to find the ground state of the spin glass at different system sizes. We could similarly observe and extrapolate the effect of the range of temperatures on the heuristic.

Future work includes using this heuristic analysis method to analyze other optimization heuristics and other problems to further verify its validity.

References

[1] M. Baity-Jesi, R. A. Baños, A. Cruz, L. A. Fernandez, J. M. Gil-Narvion, A. Gordillo-Guerrero, D. Iñiguez, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, J. Monforte-Garcia, A. M. n. Sudupe, D. Navarro, G. Parisi, S. Perez-Gaviro, M. Pivanti, F. Ricci-Tersenghi, J. J. Ruiz-Lorenzo, S. F. Schifano, B. Seoane, A. Tarancon, R. Tripiccione, and D. Yllanes. Critical parameters of the three-dimensional ising spin glass. *Phys. Rev. B*, 88:224416, Dec 2013.

[2] F. Barahona. On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241–3253, oct 1982.

[3] A. L. Blum and R. L. Rivest. Training a 3-node neural network is np-complete. *Neural Networks*, 5(1):117–127, 1992.

[4] S. Boettcher. Doe-ascr pre-proposal: Performance measures for optimization heuristics.

[5] S. Boettcher. Analysis of the relation between quadratic unconstrained binary optimization and the spin-glass ground-state problem. *Phys. Rev. Research*, 1:033142, Dec 2019.

[6] S. Boettcher and A. G. Percus. Optimization with extremal dynamics. *Phys. Rev. Lett.*, 86:5211–5214, Jun 2001.

[7] K. Hukushima and K. Nemoto. Exchange monte carlo method and application to spin glass simulations. *Journal of the Physical Society of Japan*, 65(6):1604–1608, 1996.

[8] K. Jonason, E. Vincent, J. Hammann, J. P. Bouchaud, and P. Nordblad. Memory and chaos effects in spin glasses. *Phys. Rev. Lett.*, 81:3243–3246, Oct 1998.

[9] H. G. Katzgraber. Introduction to monte carlo methods. arXiv:0905.1629, 2009.

[10] N. Metropolis, A. W. Rosenbluth, M. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[11] F. Zamponi. Mean field theory of spin glasses. arXiv:1008.4844, 2010.