

Quantum Physics Inspired Methods for Two-Batch PWR Loading Pattern Optimization

Joseph Fustero¹, Frank Mueller¹, David Kropaczek², Colleen Farrelly³

¹North Carolina State University, jgfuster@ncsu.edu, mueller@cs.ncsu.edu

²Oak Ridge National Laboratory, kropaczekdj@ornl.gov

³Staticsym, LLC, cfarrelly@med.miami.edu

[leave space for DOI, which will be inserted by ANS]

INTRODUCTION

In this work, a simplified pressurized water reactor (PWR) fuel loading pattern optimization problem is solved. This is a modification of a problem [1] previously studied using quantum annealing (QA), involving the design of the APR1400 initial core with quarter-core (rotational only) symmetry and the choice of two fuel assembly designs, a high and low enrichment, for each core location. Lastly, a lower-enrichment fuel assembly is fixed at the center of the core. Thus, the problem has 60 degrees of freedom and 2^{60} potential solutions. This differs from the previous study [1], which sought a 30/30 split between higher- and lower-enrichment fuel assembly design usage. Here, the constraint on region split has been removed and additional constraints were added to the cost function. The previous study considered only the maximum local power peaking factor as well as the cycle length. Here, the initial critical boron, as well as the fuel cost, are included in the cost function. Note that the cycle length and fuel cost serve to drive the solution to the optimal region split between the more expensive high enriched fuel and the low enriched fuel. Candidate loading patterns are evaluated by the metrics listed in Table I.

TABLE I. Cost Function Components

Component	Limit	Penalty	Credit
Initial Boron (B, ppm)	1450	10	0
Maximum Local Peaking Factor (F_q)	2.15	100	10
Cycle Length (CE, $\frac{GW-d}{MT}$)	20.95	100	0
Fuel Cost (FC, $\frac{\$}{MW-h}$)	10.06	100	100

With H as the Heaviside step function, the cost function for a given loading pattern is given by Eq. (1).

$$E = 10H(B - 1450) + 100H(F_q - 2.15) - 10H(2.15 - F_q) + 100H(CE - 20.95) + 100H(FC - 10.06) - 10 - 0H(10.06 - FC), \quad (1)$$

In Eq. (1), B represents the initial boron, F_q represents the maximum local power peaking factor, CE represents the cycle length, and FC represents the fuel cost. Acceptable loading patterns must meet the upper limits on the initial boron, maximum local power peaking factor and fuel cost and the lower limit on the cycle length.

METHODS

The phrase “quantum computing” generally refers to (i) adiabatic quantum computing (ii) gate-based quantum computing or (iii) the implementation of algorithms inspired by quantum physics on a traditional (classical) computer. Here, only (i) and (iii) are addressed. Addressing (i), the D-Wave 2000Q was used to solve the problem. Additionally, a quantum-inspired genetic algorithm (QGA) was used to solve the problem, addressing (iii). QGA is a variation on more typical genetic algorithms (GAs) that have been used in solving fuel loading pattern optimization problems. None of the gate-based architectures (ii) were used in solving this problem. Lastly, parallel simulated annealing (PSA) based on the constraint annealing method was used to solve the stated problem [2].

The CASMO-SIMULATE code system for reactor analysis was used to evaluate the cost components of the candidate loading patterns. This involves (1) one-time generation of the nuclear cross-section library using CASMO for the fuel assembly designs in the problem and then (2) core depletion using SIMULATE to determine the metrics listed in Table I. Here, PSA and QGA could be “traditional” methods in that they generate sample loading patterns as input to SIMULATE, and then, based upon the SIMULATE output, they adaptively generate new samples. Simulated annealing (SA) and classic genetic algorithms (GAs) are also common methods but these results are not presented here [3].

Quantum computers, in general, are useful for solving fewer problems than what is commonly believed. D-Wave solves the quadratic unconstrained binary optimization (QUBO) problem. Other methods exist for solving QUBO problems, and the authors make no statement on whether D-Wave is superior or could be made superior to these other methods [4]. Nevertheless, an examination of how to relate the fuel loading optimization problem to a QUBO problem

will allow insight into the fuel loading optimization problem. Additionally, D-Wave solutions to the QUBO representing the problem must be input into SIMULATE and the output must be extracted to determine solution quality. D-Wave, then, unlike the “traditional” methods, does not work by directly manipulating the input and output of SIMULATE. If the best solutions generated by D-Wave are comparable to the best solutions generated by PSA, then the method by which D-Wave solves the fuel loading problem (i.e. by solving a QUBO) should be examined in greater detail. QUBOs are a specific type of sparse graph and-if shown to be useful in modeling the fuel loading problem-should be examined in greater detail as surrogate models for the fuel loading problem.

A process [1] for using CASMO-SIMULATE to generate a QUBO was previously developed assuming random sampling (20,000 samples) of SIMULATE. This process creates a surrogate model using a machine learning approach; however, other methods have directly created the QUBO by embedding a rules-based approach [5]. First, since the neutron mean free path in PWRs is small, it is appropriate to use a 2D Ising model (easily convertible to a QUBO) that only incorporates nearest-neighbor connectivity between fuel assemblies. Once the graph’s structure is assumed, the values for its vertices (linear coefficients) and edges (quadratic coefficients) must be trained. First, acceptable fuel loading patterns were sampled. Secondly, the samples were input to SIMULATE, and the cost functions were evaluated according to Eq. (1). Next, the ground state of the cost function is estimated. The lower-enrichment fuel assemblies were assumed to have a spin of “-1” whereas the higher-enrichment fuel assemblies were assumed to have a spin of “1.” The sign convention on the summation terms is the same as that of the previous methodology. A previously discovered training algorithm [6] for the 2D Ising model was altered to include corrections to account for the expected Boltzmann distribution of the solutions. Additionally, the tuning parameter was set to zero, as the 2D Ising model is not a very accurate surrogate model (though it is more accurate at lower energies). The algorithm was run and the values for the linear and quadratic coefficients of the 2D Ising model were discovered. For the current problem, the only alteration is that the sampling randomly selects (with equal probability) a fuel assembly design at each location as opposed to randomly selecting an entire fuel loading pattern at once with a 30/30 split in the fuel assembly designs. The 2D Ising model obtained using this method is shown in Fig. 1. In all 2D Ising maps (Figs.1, 2, and 3), the blue locations represent the location of a fuel assembly as well the corresponding linear coefficient to the 2D Ising model. Additionally, the green locations represent the quadratic coefficient, the coupling between fuel assembly locations. Before being input into D-Wave, all 2D Ising models were converted to QUBOs.

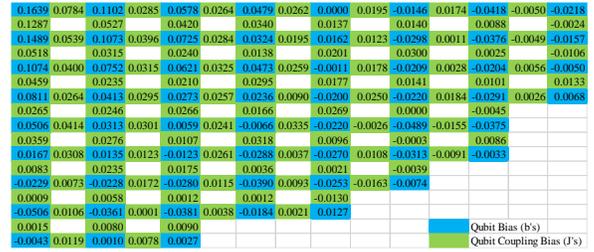


Fig. 1. 2D Ising map generated using random samples.

As seen in Fig. 1, the behavior of the linear and quadratic coefficients is broadly as expected. The quadratic coefficients are largely positive, indicating a preference for dissimilar fuel assembly designs to be located adjacent to each other. This is consistent with checkerboard-style strategies. Additionally, the linear coefficients begin as more positive towards the center of the core (favor lower-enrichment fuel) but become more negative closer to the periphery (favor higher-enrichment fuel) of the core. This behavior is consistent with IN-OUT fuel loading pattern strategies.

Next, two additional QUBO models were prepared. Parallel simulated annealing was run to generate approximately 25,000 candidate loading patterns. For the first 20,000 of these loading patterns, the set comprising those meeting the criteria listed in Table I was found. Thus, all of the first 20,000 solutions as well as only the solutions meeting the criteria in Table I among the first 20,000 solutions were used to train a 2D Ising model. The major change from the method used to train the 2D Ising model for the random cases is that there is no longer the exponential correction term accounting for the Boltzmann distribution. Because PSA itself is an optimization method, it should favor the lower energy solutions. Fig. 2 shows the 2D Ising model trained on the first 20,000 solutions generated by PSA and Fig. 3 shows the 2D Ising model trained on only accepted solutions among the first 20,000 solutions generated by PSA.



Fig. 2. 2D Ising map generated using all PSA samples.

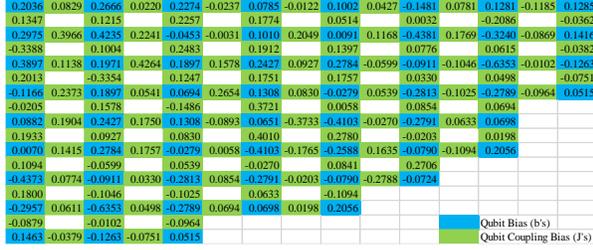


Fig. 3. 2D Ising map generated using only accepted PSA samples.

Comparing Fig. 2 and Fig. 3 to Fig. 1, it is apparent that for the PSA-informed 2D Ising models, the outermost linear coefficients (unlike in Fig. 1) are negative, implying a preference for the lower-enriched fuel assemblies at this location. Thus, unlike the 2D Ising model trained on random cases, the 2D Ising models presented in Figs. 2 and 3 capture the low neutron leakage design criterion. Otherwise, the 2D Ising models in Figs. 2 and 3 are broadly similar both to each other as well as those in Fig. 1. All models point to a preference for checkerboard-style loading patterns as well as IN-OUT loading pattern strategies. To determine which model best captures the details of the problem, the QUBOs (2D Isings are easy to convert to QUBOs) are embedded onto D-Wave to solve the QUBO using quantum annealing. To keep the results of all methods comparable, at most 5,000 unique solutions should be generated using D-Wave, with the solutions subsequently input to and evaluated with SIMULATE.

Likewise, QGA should be used to evaluate 25,000 candidate solutions for a valid comparison of computational expense. QGAs implement the usual genetic operators of randomly initializing a chromosome population (here, qubits rather than bits), evaluating the fitness of chromosomes, and mutating chromosomes to generate new populations. QGAs add a qubit rotation gate that informs chromosome mutation, allowing for superposition to hasten convergence. Prior results suggest that QGAs converge faster and find solutions of similar quality compared to GAs on loading pattern optimization problems [7]. QGA solutions provide similar best loading patterns to PSA solutions.

RESULTS

Table II lists the five methods along with the evaluated cost function for the best loading pattern generated. Table II shows that using QA to solve the 2D Ising trained only on accepted PSA samples has a much higher evaluated cost function (worse solution) than the other four methods, which obtain best solutions of comparable quality. Figs. 4, 5, 6, 7, and 8 show graphically the best loading pattern obtained using each method. In these figures, blue represents lower-enrichment fuel assemblies and red represents the higher-enrichment fuel assemblies.

TABLE II. Evaluated Cost Function

Method	Average	Best
PSA	171.2	-0.5
QGA	513.3	31.4
QA, Random	180.7	57.4
QA, All PSA	177.7	34.3
QA, Accepted PSA	348.9	212.0

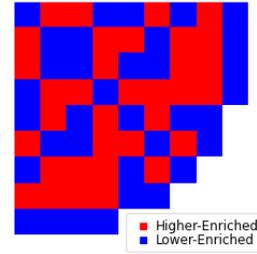


Fig. 4. Best loading pattern generated by PSA.

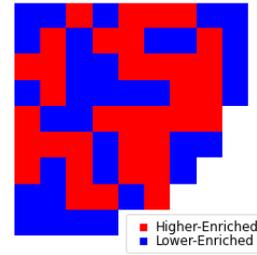


Fig. 5. Best loading pattern generated by QGA.

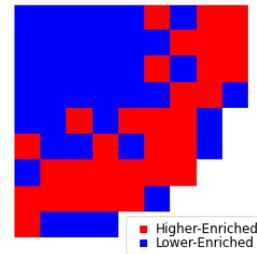


Fig. 6. Best loading pattern generated by QA solve of 2D Ising trained on random samples.

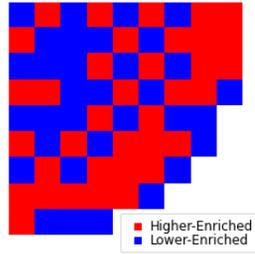


Fig. 7. Best loading pattern generated by QA solve of 2D Ising trained on all PSA samples.

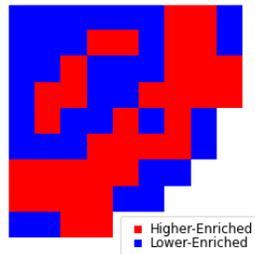


Fig 8. Best loading pattern generated by QA solve of 2D Ising trained on only accepted PSA samples.

In general, the best loading patterns generated using QA (Figs. 6, 7, and 8) have much more higher-enrichment fuel on their periphery compared to the loading patterns generated using methods unrelated to solving a 2D Ising surrogate (Fig. 4: PSA, Fig. 5: QGA). Although the result in Fig. 7 has nearly the same cost function evaluation as that in Fig. 5, the Fig. 5 result is most similar to the Fig. 4 result. In fact, the PSA and QGA results are remarkably similar in terms of the structure of the loading pattern closer to the core. On the other hand QA results either mimic IN-OUT strategies (Fig. 8) or IN-OUT strategies combined with checkerboard-style loading patterns (Figs. 6 and 7).

CONCLUSIONS

PSA, QGA and QA were used to solve a two-batch PWR fuel loading optimization problem. PSA and QGA resulted in best solutions of very similar quality. When QA was used to solve 2D Ising surrogate models trained on either random samples or all PSA results, the best QA solution was comparable to the best PSA and QGA solutions. However, when QA was used to solve the 2D Ising trained on only PSA samples that met the design criteria, the best solution was much worse than any of the other four approaches taken. A surprising result is that for QA approaches, both using samples resultant of some classical optimization method (such as PSA) as well as explicit reconstruction of the Boltzmann distribution result in the same solution quality. Failure to include lower quality solutions during the classical

optimization approach also results in a lower quality best solution obtained using QA.

It must be re-stated that the results in this paper do not reflect conclusions on the validity of quantum computing itself; rather, this work explores how the steps involved in formulating the problem into a QUBO allows for greater insight into the problem. QA is a competitive option for generating loading patterns despite the low surrogate accuracy [1] of the 2D Ising model. Thus, the general use of sparse graphs as surrogate models for fuel loading pattern optimization should be investigated further for reactors with tight local spatial coupling. Further innovation could be including diagonal coupling to improve the sparse graph surrogate accuracy.

ACKNOWLEDGMENTS

This work was supported in part by the following awards: NSF 1917383, 1818914, 2120757, and LANL subcontract 725530.

REFERENCES

1. J. FUSTERO, S. PALMTAG and F. MUELLER, "Quantum Annealing Stencils With Applications to Fuel Loading of a Nuclear Reactor." *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pp. 265-275, (2021). doi: 10.1109/QCE52317.2021.00044.
2. D. KROPACZEK and R. WALDEN, "Constraint Annealing Method for Solution of Multiconstrained Nuclear Fuel Cycle Optimization Problems," *J. Nuclear Science and Engineering*, 193, 506-522 (2019).
3. B.D. ANDERSON, "A Machine Learning Based Approach to Minimize Crud Induced Effects in Pressurized Water Reactors." Dissertation, North Carolina State University (2021).
4. J. KING, ET. AL., "Quantum Annealing Amid Local Ruggedness and Global Frustration." *Journal of the Physical Society of Japan*, **88**, 061007, (2019) 10.7566/JPSJ.88.061007
5. A. WHYTE and G. PARKS, "Quantum Annealing Optimization of a Heuristic Surrogate Model for PWR Fuel Loading." *Proc. PHYSOR 2020*, Cambridge, United Kingdom, March 29–April 2, 2020, American Nuclear Society (2020).
6. A. LOKHOV, M. D. VUFFRAY, S. MISRA & M. CHERTKOV, "Optimal Structure and Parameter Learning of Ising Models." LA-UR-16-294425, Los Alamos National Laboratory, (2012).
7. A. DOS SANTOS NICOLAU, R. SCHIRRU, & A. M. M. DE LIMA, "Nuclear Reactor Reload Using Quantum Inspired Algorithm." *Progress in Nuclear Energy*, **55**, 40-48, (2012).