A Hamiltonian Engine for Radiotherapy Optimization

Mehrdad Malekmohammadi[1], Keivan Dabiri[1],Joshua Mathews[2],
Daryl P. Nazareth[2], Hirotaka Tamura[3], and Ali Sheikholeslami[1]

***Abstract*— We apply a new hardware and software platform called the Hamiltonian Engine for Radiotherapy Optimization (HERO) to the problem of Intensity-Modulated Radiation Therapy (IMRT) treatment planning. HERO solves large general-form binary optimization problems by decomposing them into sub-problems and approximating them using a quadratic pseudo-boolean function. Optimizing the resulting function becomes a quadratic unconstrained binary optimization (QUBO) problem, which has been widely studied and has numerous applications in various fields. A Quantum Annealer (QA) approach has been previously investigated to solve QUBO problems, including IMRT optimization. However, the QA can only accommodate a small number of variables and requires several hours to obtain optimized plans. HERO acts as an optimizer for QUBO problems, which not only addresses these shortcomings but also relies solely on conventional hardware design while operating at room temperature. We evaluate HERO on seven prostate IMRT cases with clinical objectives, each using approximately 6000 beamlets. Our method was compared to the commercial treatment planning software, Eclipse, for both time-to-solution and plan quality. HERO solves most cases in about 30 seconds, with significantly lower objective function scores than Eclipse. The results indicate that HERO is promising for radiation therapy optimization problems. Additionally, HERO has the potential to be applied to Volumetric-Modulated Arc Therapy (VMAT) and other complex types of treatment planning.**

Keywords: Radiation therapy, Optimization, Hardware

I. INTRODUCTION

Intensity-Modulated Radiation Therapy (IMRT) involves patient irradiation using a fixed beam arrangement and dynamic fluence fields created by a multileaf collimator. The spatially-varying intensities of the fields are optimized according to certain objectives, typically requiring the prescribed dose to be delivered to the tumor, while maintaining the doses to adjacent organs at risk (OARs) below clinical tolerances. The inverse planning methods required for IMRT treatments comprise a mature field of research, and various algorithms have been applied to this problem. These algorithms attempt to minimize an objective function defined based on certain clinical criteria.

Many of the algorithms employed in commercial systems for IMRT optimization are gradient-based and offer fast performance, but are prone to converging to non-optimal solutions [1]. It has been shown that the objective functions based on dose-volume objectives consist of several local minima due to their non-convexity [2]. However, some argue that in many cases, local minima can be avoided if an effective starting point is chosen or a convex objective function is used [1], [3]. To this end, several works use linear programming models in beamlet intensity optimization [4], [5], [6]. Regardless, managing the dose-volume constraints in a convex formulation can be challenging.

Stochastic search methods like simulated annealing have the capability of escaping from local minima due to their intrinsic random characteristics, but are generally slow [7], and therefore not usually employed clinically.

In addition to conventional algorithms, a recent work [8] has investigated the use of Quantum Annealing (QA) in IMRT beamlet optimization. QA is a combined hardware and software platform for solving combinatorial optimization problems, developed by D-Wave (D-Wave Inc., Burnaby, BC). This approach required reformatting the IMRT problem into a form suitable for the QA approach. Similar to simulated annealing, though implemented in actual hardware, QA attempts to escape local minima in the objective function by exploiting quantum tunneling effects at near absolute zero temperatures [9]. Despite its capability for solving certain combinatorial problems rapidly, the QA required more than three hours to optimize a prostate case with only 70 beamlets, and the quality of resulted plans were not clinically acceptable [8].

It has been shown that simulated annealing, when combined with the massive parallelism offered by certain hardware platforms, can lead to novel hardware architectures that provide significant speedup compared to the standard simulated annealing algorithm [10], [11].

Inspired by the prior work in [11] and the concept of Markov Chain Monte Carlo (MCMC) search, we propose and implement a new heterogeneous hardware platform, called the Hamiltonian Engine for Radiotherapy Optimization (HERO), which can be applied to complex optimization problems in radiation therapy, such as IMRT and VMAT. This platform addresses both issues of handling large numbers of variables and casting the problem in a format supported by the engine. In this research, we apply HERO to IMRT beamlet optimization for prostate cancer cases and compare the results to that of a commercial treatment planning system. This work also includes the study of different objective functions in terms of plan quality and time complexity, as well as a GPU implementation for large-scale parallelization of objective function calculations.

[1]The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto

[2]Department of Radiation Medicine, Roswell Park Comprehensive Cancer Center, Medical Physics, University at Buffalo

[3]Fujitsu Laboratories Ltd., Kawasaki, Japan

II. MATERIALS AND METHODS

A. Objective Function

We investigated the suitability of different objective functions for the IMRT beamlet optimization of prostate cases. In each case, five equally-spaced fields were arranged in the Computational Environment for Radiotherapy Research [12] (CERR). The dose for each field was calculated and decomposed into square beamlets of approximately 0.25 cm per side.

The purpose of the optimization process is to ensure that the target volume receives the prescribed dose, while the OARs are prevented from receiving doses beyond certain objectives. The objective function is therefore formed by penalizing the Dose-Volume Histogram (DVH) of each structure with respect to its dose-volume objective. These clinical objectives are usually in the form of either overdose or underdose specifications. We aim to minimize an objective function that consists of a squared sum of penalty terms over all the objectives.

B. MCMC Optimizer

The MCMC optimizer utilizes stochastic search methods to solve combinatorial optimization problems, specifically the Quadratic Unconstrained Binary Optimization (QUBO) problem. The objective here is to minimize a quadratic pseudo-boolean function in the following format:

$$f(\delta_1, \delta_2, ..., \delta_n) = -\sum_{i<j} W_{ij}\delta_i\delta_j - \sum_i b_i\delta_i, \tag{1}$$

where δ_i's are binary variables taking values in $\{0,1\}$, W_{ij} represents the connection strength between variables δ_i and δ_j, and b_i is the bias term for variable δ_i.

Using stochastic search methods and parallel computing, the MCMC optimizer is able to minimize an energy function in the form of Eq. 1. The stochastic search performed by the MCMC optimizer enables it to escape possible local minima in an energy landscape.

C. Hamiltonian Engine for Radiotherapy Optimization (HERO)

In recent years, the concept of simulated annealing, combined with massive hardware parallelism, has been proven effective in solving optimization problems in many applications [10], [13], such as in transportation (traveling salesman problem) and physics (spin-glass problems). Indeed, this approach has led to a hardware platform that can solve QUBO problems with up to 1024 fully-connected variables [10]. However, most real-world problems have a larger scale and may not be readily available in QUBO form. In this work, we address these two issues by implementing a surrogate-based optimization algorithm and adapting it to IMRT optimization problems. In order to deal with the large number of optimization variables, we partition the large problem into smaller pieces, and apply the MCMC optimizer to solve these pieces. Moreover, we propose a method to approximate any general combinatorial optimization problem by a quadratic function for the purpose of optimization.

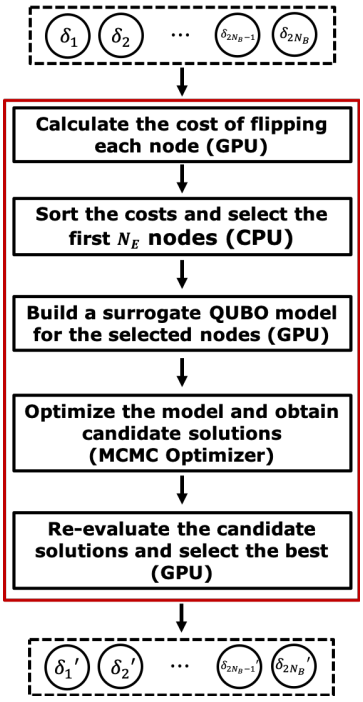

Fig. 1. Block diagram illustrating the main algorithm of HERO. The red block includes the main steps in the HERO algorithm and the dashed blocks represent the input and the output of the algorithm.

1) Input Variables: We control the value of each real variable by using binary variables. Assume we have a set of beamlets with intensities $\boldsymbol{x} = \{x_1, x_2, ..., x_{N_B}\}$, which are non-negative real variables. A set of binary nodes $\boldsymbol{\delta} = \{\delta_1, \delta_2, ..., \delta_{2N_B}\}$ are defined such that each two nodes correspond to one beamlet, and will determine whether the intensity values will be increased or decreased by a certain *step* value, or remain the same. The step size is selected as 25% of the initial value of each beamlet weight. The binary nodes become the binary optimization variables, and are updated based on the algorithm described below. This process is repeated for each iteration with a new configuration of beamlet intensity values, which has been modified based on the previous iteration.

2) Main Algorithm: The high-level algorithm as illustrated in Fig. 1 consists of five main steps, with the hardware employed for each step indicated. The input to the algorithm is a vector of the binary values associated with each node, defined for the current intensity value of each beamlet. The δ_i's are initially all set to zero, indicating no initial change to the current intensity values. In the first step, we calculate the cost of flipping each node, defined as the change in the objective function value if a node is flipped from 0 to 1, and thus calculate the cost for $2N_B$ nodes, where N_B is the total number of beamlets. In the next step, we sort the costs from lowest to highest, and select a portion of the nodes to be submitted to the MCMC optimizer. At this point, we build a surrogate QUBO model over the selected nodes, and then optimize the model by running the MCMC optimizer for a certain amount of time. This provides candidate solutions,

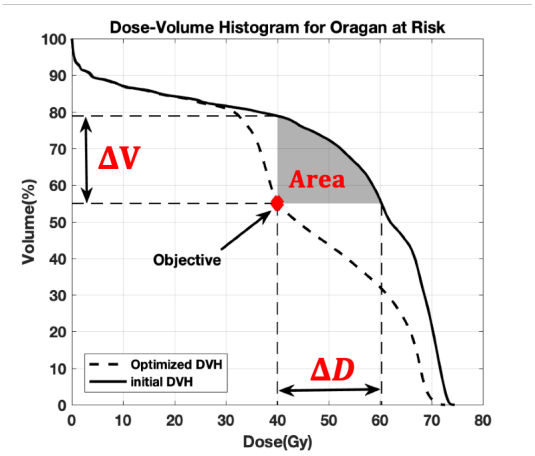

Fig. 2. An illustration of an OAR DVH being penalized based on dose difference vs volume difference. The solid-line curve shows the optimized DVH in which both terms are zero.

which are evaluated to identify the solution with the lowest objective function value. The resulting binary code is used to update the real intensity values, and the entire process is repeated.

3) Cost Calculation and Partitioning: To calculate the objective function, the amount of overdose/underdose is computed for all objectives. Fig. 2 illustrates an example of a DVH graph with one objective for which the ΔD must be evaluated. This requires calculating DVHs for each structure using a sorting process applied to all the voxel doses. Since the structure may consist of $100,000$ or more voxels, sorting becomes computationally intensive. We therefore attempt to exploit the parallelism of GPUs to improve the computational efficiency of cost calculation. We propose a GPU implementation to calculate the ΔD using a binary search-like method, by iteratively calculating the ΔV for a certain dose and adjusting the dose respectively until ΔV becomes zero.

To calculate the ΔV, we multiply the dose deposition matrix ($\boldsymbol{D}$) by the intensity vector ($\boldsymbol{x}$) to obtain the total dose received by each voxel. Subsequently, we compare the dose vector with the relevant dose objective(s) and count the number of voxels with larger dose. Although this approach has a similar time complexity to sorting, the performance of the implementation in GPUs is faster.

In order to select a portion of the variables to be submitted to the MCMC optimizer, we sort all the nodes by their costs, using a quicksort algorithm, and then select N_E nodes with the lowest costs.

4) Surrogate Model: Using the N_E nodes selected in the previous section, we build a surrogate QUBO model while fixing all nodes outside the subset. The model is based on the information in the immediate neighborhood, defined by all the states within the hamming distance of 2 bits or less from the current solution (the hamming distance is the number of differing bits in two strings of bits). As shown in Eqs. 2-4, this information is obtained by flipping 2 distinct nodes from 0 to 1, δ_i and δ_j, and computing the change in the objective function, $\Delta f_{\delta_i,\delta_j}$.

$$f(\delta_1,\delta_2,...,\delta_{N_E}) = -\sum_{\{i,j\}} W_{ij}\delta_i\delta_j - \sum_i b_i\delta_i, \tag{2}$$

$$\Delta f_{\delta_i,\delta_j} = -W_{ij} - b_i - b_j, \tag{3}$$

$$W_{ij} = -\Delta f_{\delta_i,\delta_j} - b_i - b_j. \tag{4}$$

By design, this model has zero error when the hamming distance is 1 or 2, but has increasing error for higher hamming distances. To reduce this error, instead of recording only the minimum solution found by the MCMC optimizer, we re-evaluate many candidate solutions and select the solution with the lowest objective function value.

D. Evaluations

We evaluated the optimization method retrospectively on seven prostate cases, with actual clinical CT and structure data, and the treatment planning objectives based on an RTOG protocol. We compared HERO to the commercial treatment planning system Eclipse, version 15.6 (Varian Inc., Palo Alto, CA), which used the Analytical Anisotropic Algorithm for dose calculation, and the dose-volume objectives for optimization. To apply HERO to each plan, we used CERR to arrange 5 equally-spaced beams and calculate pencil beams of approximately $0.25\ cm \times 0.25\ cm$, using the Quadrant Infinite Beam (QIB) dose calculation algorithm. This resulted in approximately 6000 beamlet dose matrices for each case. In every case, the dose voxel size used was $1.3\ mm \times 1.3\ mm \times 2.5\ mm$, equal to the CT voxel size. To calculate the costs more quickly, we subsampled the voxels and the beamlets with the option to progressively increase the resolution if needed. We used a subsampling factor of 4, which does not significantly reduce the accuracy of the objective function calculation. The size of the MCMC optimizer was set at 50 (though any value less than 1024 could be selected).

The HERO halts if a score drops below 0.01% of its initial value, or after 5 minutes. The Eclipse stopping condition was set to the default criteria of 100 minutes or 1000 iterations. However, during the optimization process, Eclipse's graphical progress indicator was monitored. The elapsed time was measured from the start until the last improvement was made.

III. RESULTS

A. Treatment Plan Optimization

Table I includes the dose values for each DVH objective, averaged over the seven prostate cases. The doses received by each organ after optimization by HERO and Eclipse are indicated. The HERO method produces the lowest OAR doses, while still providing satisfactory coverage of the PTV.

In Fig. 3 the DVHs for one of the patients are shown. The HERO results in higher quality DVHs for the OARs, while Eclipse provides slightly more homogeneous DVHs for the PTV. Fig. 4 includes a log-log plot that shows the scores versus time to best solution for the HERO and Eclipse plans. Both HERO and Eclipse optimized all plans in less

TABLE I

COMPARISON OF DOSE RECEIVED BY EACH ORGAN AVERAGED OVER SEVEN PATIENTS USING TWO DIFFERENT APPROACHES

Structure (Objective)	Eclipse	HERO
PTV (95%, 66.6 Gy)	66.6 ± 0.0	66.6 ± 0.0
PTV (max, 73 Gy)	70.4 ± 0.3	72.7 ± 0.7
PTV (min, 63.3 Gy)	61.7 ± 0.6	62.7 ± 0.5
Bladder (70%, 40 Gy)	19.6 ± 13.8	17.6 ± 13.9
Bladder (50%, 65 Gy)	34.3 ± 14.5	33.3 ± 15.0
Rectum (55%, 60 Gy)	46.7 ± 16.6	35.5 ± 13.4
Rectum (35%, 65 Gy)	63.3 ± 2.9	60.2 ± 3.2
LFH (10%, 50 Gy)	33.1 ± 3.3	29.6 ± 3.6
RFH (10%, 50 Gy)	32.8 ± 2.7	29.3 ± 2.5

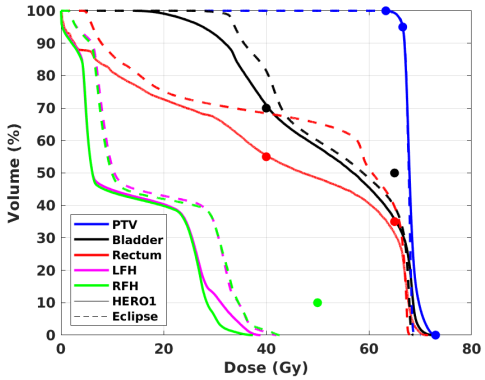

Fig. 3. Comparison of DVHs for one of the patients optimized using Eclipse and HERO. Objectives are shown by filled circles with same color as their corresponding organ.

than 60 seconds. However, HERO provided a lower objective function for every plan.

Note that the doses for both the Eclipse and HERO plans are calculated before leaf-sequencing. The beamlet intensities produced by the HERO method ranged from 0.0 to 1.7, with steps of 0.25. Compared to Eclipse, where the intensities have real values, the low number of levels of HERO may provide an advantage during the leaf sequencing process.

IV. CONCLUSIONS

We have developed a heterogeneous computational platform, consisting of a CPU, GPU, and an MCMC optimizer. By building a QUBO model of the IMRT optimization problem, we have applied our method to beamlet-weight optimization using clinical DVH objectives. We compared the method to the commercial clinical software, Eclipse, by applying both systems to seven prostate cases, demonstrating that the HERO method produced higher-quality treatment plans that were optimized within one minute, and in some cases, in less than 30 seconds. This new hardware-software platform may be scaled to VMAT problems in the future.

V. ACKNOWLEDGMENTS

We would like to thank Shawn Mattot, Senior Systems Analyst at Roswell Park Comprehensive Cancer Center; Matthew Jones of the Center for Computational Research, University at Buffalo; and Mohammad Bagherbeik from University of Toronto, for their valuable suggestions. We would also like to thank Fujitsu Laboratories Ltd. and Natural Sciences and Engineering Research Council of Canada (NSERC) for funding this research.

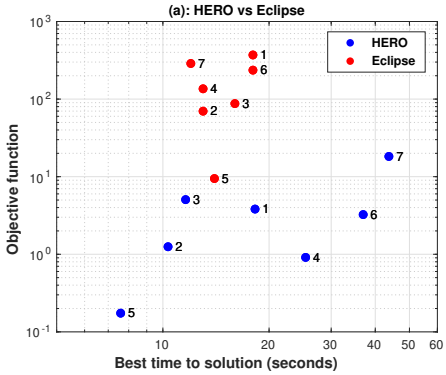

Fig. 4. A log-log plot of scores and execution time of 7 patients optimized by HERO and Eclipse. The points for HERO are an average of 10 runs.

REFERENCES

[1] X. Zhang, H. Liu, X. Wang, L. Dong, Q. Wu, and R. Mohan, "Speed and convergence properties of gradient algorithms for optimization of imrt," *Medical physics*, vol. 31, no. 5, pp. 1141–1152, 2004.

[2] Q. Wu and R. Mohan, "Multiple local minima in imrt optimization based on dose–volume criteria," *Medical physics*, vol. 29, no. 7, pp. 1514–1527, 2002.

[3] M. Alber, G. Meedt, F. Nüsslin, and R. Reemtsen, "On the degeneracy of the imrt optimization problem," *Medical Physics*, vol. 29, no. 11, pp. 2584–2589, 2002.

[4] S. Breedveld, B. van den Berg, and B. Heijmen, "An interior-point implementation developed and tuned for radiation therapy treatment planning," *Computational Optimization and Applications*, vol. 68, no. 2, pp. 209–242, 2017.

[5] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, and A. Kumar, "A new linear programming approach to radiation therapy treatment planning problems," *Operations Research*, vol. 54, no. 2, pp. 201–216, 2006.

[6] H. H. Zhang, R. R. Meyer, J. Wu, S. A. Naqvi, L. Shi, and W. D D'Souza, "A two-stage sequential linear programming approach to imrt dose optimization," *Physics in Medicine & Biology*, vol. 55, no. 3, p. 883, 2010.

[7] T. Bortfeld, "Optimized planning using physical objectives and constraints," in *Seminars in Radiation Oncology*, vol. 9, no. 1. Elsevier, 1999, pp. 20–34.

[8] D. P. Nazareth and J. D. Spaans, "First application of quantum annealing to imrt beamlet intensity optimization," *Physics in Medicine & Biology*, vol. 60, no. 10, p. 4137, 2015.

[9] M. W. Johnson, M. H. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk *et al.*, "Quantum annealing with manufactured spins," *Nature*, vol. 473, no. 7346, p. 194, 2011.

[10] S. Matsubara, H. Tamura, M. Takatsu, D. Yoo, B. Vatankhahghadim, H. Yamasaki, T. Miyazawa, S. Tsukamoto, Y. Watanabe, K. Takemoto *et al.*, "Ising-model optimizer with parallel-trial bit-sieve engine," in *Conference on Complex, Intelligent, and Software Intensive Systems*. Springer, 2017, pp. 432–438.

[11] K. Dabiri, M. Malekmohammadi, A. Sheikholeslami, and H. Tamura, "Replica exchange mcmc hardware with automatic temperature selection and parallel trial," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 7, pp. 1681–1692, 2020.

[12] J. O. Deasy, A. I. Blanco, and V. H. Clark, "Cerr: a computational environment for radiotherapy research," *Medical physics*, vol. 30, no. 5, pp. 979–985, 2003.

[13] M. Aramon, G. Rosenberg, T. Miyazawa, H. Tamura, and H. G. Katzgraber, "Physics-inspired optimization for constraint-satisfaction problems using a digital annealer," *arXiv preprint arXiv:1806.08815*, 2018.