Optimization of Evolutionary Strategy using Island Model to Design HIFU Treatment Plans

Filip Kuklis
1st year, full-time study
Supervisor: Jiri Jaros
Centre of Excellence IT4Innovations, Faculty of Information Technology, Brno University of Technology
Bozetechova 1/2, 612 66 Brno, Czech Republic
ikuklis@fit.vutbr.cz

***Abstract*— High Intensity Focused Ultrasound (HIFU) is an emerging technique for non-invasive cancer treatment where malignant tissue is destroyed by thermal ablation. The treatment plan consists of a series of HIFU sonications since one sonication destroys only a small volume of tissue. The high-quality treatment plans with a precise HIFU positioning need to be generated to destroy all malignant tissue and prevent damage to surrounding healthy tissue. Here, we present an optimized evolutionary strategy that uses an island model to design HIFU treatment plans using patient-specific material properties and a realistic thermal model. Although the original strategy was able to create high-quality treatment plans, the execution time was too high. The proposed version allows finding the solution more than 6-times faster, moreover, with a higher success rate.**

Keywords*— *treatment planning; HIFU; evolutionary strategy; island model

I. Introduction

In last years, High-Intensity Focused Ultrasound (HIFU) has been used to treat a variety of solid malignant tumors in a well-defined volume, including the pancreas, liver, prostate, breast, uterine fibroids, and soft-tissue sarcomas. The main benefits of HIFU over the conventional tumor/cancer treatment modalities, such as open surgery, radio- and chemo-therapy, is its non-invasiveness. Furthermore, it is non-ionizing and has fewer complications after treatment. To this day over 100,000 cases have been treated throughout the world with great success [1].

The basic principle of thermal HIFU treatment is to raise the temperature by several tens of degrees so that the tissue is destroyed via coagulative necrosis with delivering adequate ultrasound energy to the targeted area. The HIFU beam focusing results in cytotoxic levels of temperature only at a specific location within a small volume (e.g., about 1mm in diameter and about 10 mm in length), which minimizes the potential for thermal damage to tissue outside the focal region. The boundary between disrupted cells and normal tissue is typically less than 50 μm in width [2].

Large tumors can be destroyed by producing a contiguous lesion lattice encompassing the tumor and appropriate margins of surrounding tissue. However, complications may develop if vital blood vessels adjacent to the tumors are severely damaged. Moreover, blood perfusion may carry away a significant amount of energy and deteriorate the treatment outcome [3].

Nevertheless, of the advantages of HIFU, this method still suffers from delivery precision in contrast of other established therapies such as radiotherapy. Also, the treatment takes a lot of time. With recent advances in numerical methods and high-performance computing, detailed simulations accurately capturing the relevant physical behavior of focused ultrasound waves and temperature distribution in heterogeneous tissue are now possible [4].

The ultimate goal of my dissertation thesis is to automatically design ultrasound treatment plans. To produce precise plans, the model has to consist of a high-quality ultrasound model, thermal model, and tissue model. Unfortunately, nowadays the execution times of these models are too high and finding a good plan take weeks even for a 2D case. To produce 3D plans in a reasonable time, it is necessary to use optimized or highly simplified models, but with a small impact on the quality of the solution. Further, the technique of finding the solution has to be fast and highly scalable.

This paper presents the optimization of an existing Covariance Matrix Adaptation Evolution Strategy (CMA-ES) HIFU treatment planning algorithm, which produces high-quality treatment plans in 2D without the ultrasound model. The ultrasound model was replaced with a simple heuristic affecting only the position and shape of the sonication (heating source). Yet, finding the solution still takes on average more than one day on one node of the Salomon cluster with 24 cores (2 × Intel Xeon E5-2680v3). The most time-consuming part is the fitness function evaluation, which uses a tissue realistic thermal model. The optimization of the evolutionary process is based not only on the acceleration of the thermal model simulation and rewriting the simulation to C++, but moreover, on the parallelization and optimization of the evolutionary strategy by using the island model of the evolutionary algorithm. The population on each island evolves independently. The islands communicate with each other and exchange the best individuals. This approach allows the parallelization of the evolution strategy, and also, brings some benefits of separated evolutions on islands.

II. ORIGINAL EVOLUTIONARY STRATEGY

This section describes "Design of HIFU treatment plans using an evolutionary strategy."[8], which was used in this work as the base model.

A. Covariance Matrix Adaptation Evolution Strategy

The CMA [5] describes the pairwise dependencies between variables/genes on the top of the classic ES.

In the CMA-ES, a population of λ new search points (individuals, offspring) is generated by sampling a multivariate normal distribution $\mathcal{N}(m, C)$ determined by its mean $\boldsymbol{m} \in \mathbb{R}^N$ and its symmetric and positive defined covariance matrix $\boldsymbol{C} \in \mathbb{R}^{N \times N}$. The new generation of individuals is ranked according to fitness values and then the best μ individuals are selected. The elitism is not used.

B. Solution encoding

The solution can be represented as an *n-tuple* of HIFU sonications:

$$\boldsymbol{I} = (\boldsymbol{S_1}, \boldsymbol{S_2}, \dots, \boldsymbol{S_n}),\ \boldsymbol{where}\ \boldsymbol{S_i} = (\boldsymbol{x_i}, \boldsymbol{y_i}, \boldsymbol{t_{on}}, \boldsymbol{t_{off}}), \quad (1)$$

where x and y are spatial coordinates (center of sonication) and t_{on} and t_{off} represents the duration of each sonication (heating), and the pause after each sonication (cooling).

C. Fitness function

The thermal model simulates *N* HIFU sonications applied to the medium. Areas exceeding the thermal dose threshold are identified. Correctly treated and mistreated areas are identified and their weights summed together. The resulting sum represents the fitness value. If all matter is treated correctly and there are no mistreated areas, the sum will drop to zero(global optimum).

D. Thermal model

The Pennes' bioheat equation [9] is used to compute heat diffusion:

$$\boldsymbol{\rho C \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) + W_b C_b (T - T_a) + Q} \quad (2)$$

where C and C_b are specific heat of tissue and blood $[J/kg°C]$, $W_b[kg/m^3 s]$ is a blood perfusion related parameter, T_a is the arterial temperature (assumed to be 37°C), and $Q[W/m^3]$ is the power deposited in the tissue by the ultrasound transducer.

Thermal damage is computed using the Sapareto-Dewey iso-effect thermal dose relationship [6] and then it is thresholded and killed tissue is marked.

The thermal model is implemented in Matlab using the k-Wave toolbox [7] and supports precise tissue parameter settings derived from patient-specific models of the tissue anatomy.

III. THERMAL MODEL OPTIMIZATION

More than 99% percent of the evolution time is taken by the fitness evaluation – the thermal model for heterogeneous medium. Therefore, it was decided to re-implement this model using C++, and parallelize and vectorize the code using OpenMP. The re-implemented thermal model is more than three times faster than the Matlab one, even if it is invoked directly from Matlab as a MEX function.

The scalability of the model is not ideal (see Fig. 1), mostly because of the scaling of the Fast Fourier Transforms for a given size of the medium. Because of that, we can further speed up the simulation by running multiple simulations simultaneously, each of which with fewer threads than the number of cores. For the biggest speed-up, the total number of threads must be equal to the number of processor cores, and also, the number of threads for one simulation has to be chosen based on scaling plots.

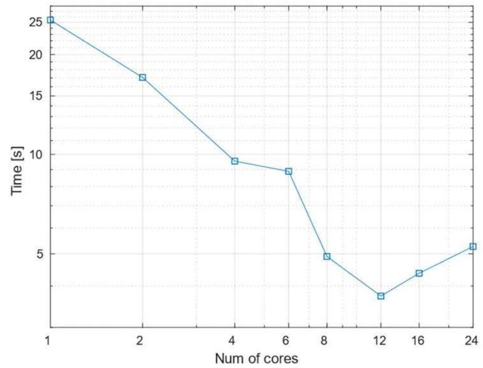

Figure 1: Strong scaling of the thermal model executed on the whole treatment plan consisting of 12 sonications.

IV. ISLAND MODEL

As mentioned in the previous section, to achieve better speed up, the thermal model has to be run in parallel. There are several ways how to do so. The simplest one is a parallel evaluation of the individuals in the population.

We chose the island model of evolution. Every island has its own independent population and the individuals migrate between different islands. Multiple migrations strategies were implemented and tested. The island model was tested on one node of the Salomon cluster and the number of islands was chosen based on the scalability of the thermal model. For the best speed-up, the condition from the previous section must be satisfied. The total number of simulations equals to the total number of cores. Therefore, the number of islands equals to the number of cores divided by threads per simulation. In Fig. 1., we can see that the model scales reasonably for 4, 8 and 12 threads which results to 6, 3 or 2 islands in our model.

A. Pseudocode

```
while isempty(stopflag)
    init_params();
    for k=1:lambda
        sp = sample_sonic_params();
        fitness(k) = eval_thermal_model(sp);
    end
    sort(fitness);
    for i=1:n_isl
        rcv = broadcast(fitness(1));
        if(rcv_cond)
            fitness(lamda-1) = rcv;
        end
    end
    update_params(fitness);
end
```

The pseudocode shows that every island communicates with each other using blocking broadcast communication. Based on the acceptance strategy, an island can or cannot accept the migrant from another island. However, communication is not dependent on that, and the islands communicate in every generation. Because of the blocking broadcast, the communication has some overhead.

The population size per island remains constant, $\lambda_{island} = 13$, so the total population size gets bigger with an increasing number of islands. If the number of islands is N, then the total population size is $\lambda_{total} = N * \lambda_{island}$.

B. Comparison with the pan-population model

The island model with different migration strategies was compared against the pan-population (PP) model with a sequentially evaluated thermal model.

To compare the island and pan-population models, the same total population size must be used. The most important parameter to investigate is the total number of evaluations. Even if the island model had the same total number of evaluations (sum of evaluations on all islands) and was faster, it would not be faster than the parallel evaluated pan-population model, because of the communication overhead. Therefore, the gain from using the island model has to be higher than the overhead caused by blocking communication.

1) 6-island model

Several migration strategies were tested on $N = 6$ islands with the total population size of $N * \lambda_{island} = 6 * 13 = 78$. In the first migration strategy (6I-SB), the island broadcasts the best individual in the current generation to other islands. In the second strategy (6I-S3), every island broadcasts 3 best individuals and the receiving island chooses one using the roulette wheel selection. In the third strategy (6I-CB), the island broadcasts 3 best individuals, but there is a probability to refuse a migrant from other island which rises with the fitness value of the best individual. The idea is to keep higher diversity between islands.

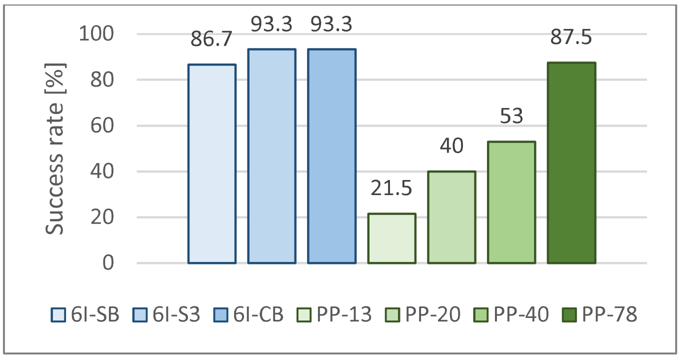

Figure 2: The success rate for the island (6 islands) and non-island model with different population sizes.

Figure 2 shows the number of evaluations to find an optimal solution. The number of evaluations of the island model is a sum of evaluations over all islands. The number of evaluations for one island is, therefore, 1/6 of the total evaluations.

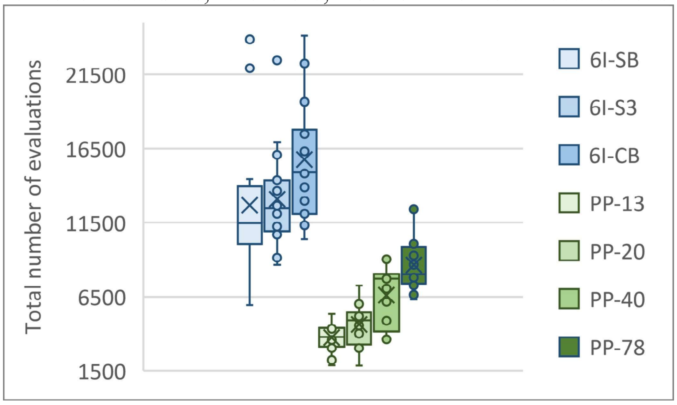

Figure 3: Number of evaluations for the island (6 islands) and pan-population model with different population sizes.

Figure 3 shows the success rate in percents of all runs which did converge to the optimum (0). Our effort is to maximize the success rate and minimize the number of evaluations to find a solution.

Figure 2 and 3 show a comparison between the island model strategies with 6 islands and PP model with 13, 20, 40 and 78 individuals in the population. We can see that the success rate of finding the solution is basically the same as the pan-population model with 78 individuals in the population, which is expected. But the total number of evaluations is significantly higher than the PP version. Therefore, the island model with 6 islands is not competitive against pan-population model with the same number of individuals in the population, even if it is faster.

2) 3-island model

Only the first migration strategy (send best individuals) was tested on 3-island model (3I-SB) with the population per island $\lambda_{island} = 13$, therefore the total population is $N * \lambda_{island} = 3 * 13 = 39$. Regarding to the population size, the comparable PP model is the one with the population size of 40 individuals.

Figure 4. shows that the 3-island model is much better at the success rate than the PP model with a population of 40 individuals, moreover, it overcomes even the PP model with 78 individuals in the population.

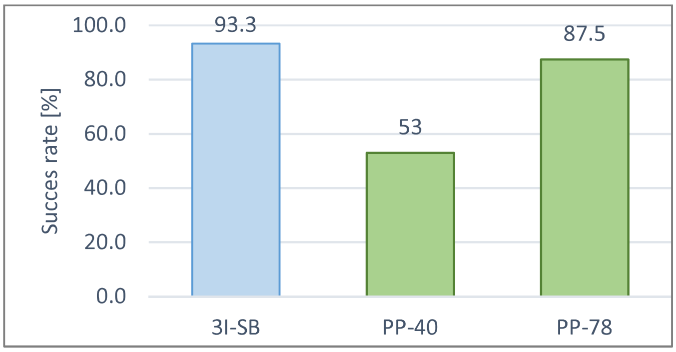

Figure 2: The success rate for the island (3 islands) and pan-population model with population size 40 and 78.

On the other hand, in the number of evaluations, the 3-island model is significantly better than the pan-population model with 78 individuals and it is even better than the pan-population model with 40 individuals in the population (see Fig. 5).

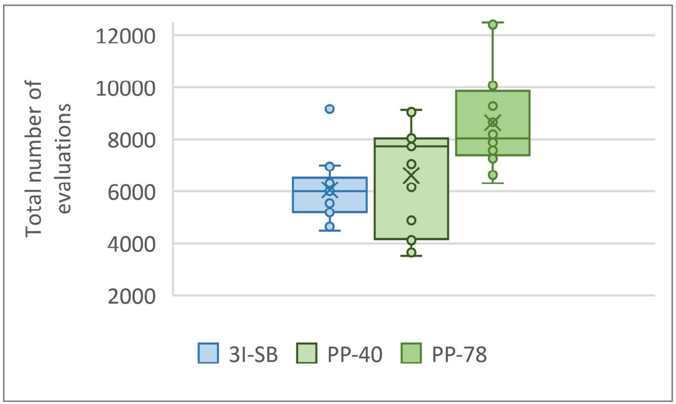

Figure 3: Number of evaluations for the island (3 islands) and pan-population model with population size 40 and 78.

The communication overhead for the 3-island model is less than 10%, therefore, the island model is much better even if the evaluation of individuals is parallelized in pan-population model.

V. Conclusion

This study showed an acceleration of the covariance matrix adaptation evolution strategy for finding the HIFU treatment plans and the relevance of using the island model.

The first optimization has been based on the reimplementation of the thermal model from Matlab to C++, and its parallelization and vectorization by OpenMP. The C++ version is implemented as a MEX function which can be called directly from Matlab. The C++ thermal model allows to run the evolutionary strategy with a much bigger population than the Matlab one. The maximum population size to find a plan in less than 48 hours with the Matlab thermal model is around 40. With the C++ thermal model, we can run the EA with more than twice as many individuals.

This study has reviewed several approaches of the island model of CMA evaluation strategy. The proposed island model is based on the idea of the parallelization of the fitness evaluation which does not scale well with the number of cores. The island model improves the algorithm by parallelization of evaluations. Moreover, the use of the island model with the less population size and number of evaluations increases the success rate over the pan-population model. Therefore, it also improves the overall time of evolution. With the combination of island model and the optimized thermal model, the evolution finds a solution more than 4 times faster on the same computer.

It must be said that not all reviewed approaches were successful and better than the comparable pan-population. The efficiency depends on the parameters of the model, for example, the number of islands.

The proposed island model also allows the parallelization of the evolutionary strategy on more interconnected computers (mainly more nodes on one cluster), which will be used in the next stages with an integrated ultrasound model.

The next most important step is to implement a simple ultrasound model and integrate it into the proposed evolution strategy. The ultrasound model will be based on a simple ultrasound ray-tracer, firstly, simulating only the attenuation, later also simulating the refractions and reflections. After the integration and testing the evolution with all models, the most critical step will be the transformation into the 3D space. The idea is to design the evolution strategy as a master-worker model, where the master runs the evolution and workers evaluate the fitness function – i.e., the model simulations will run on workers. On a higher level, the evolution will be parallelized by the proposed island model strategy.

VI. Acknowledgment

This work was supported by the FIT-S-17-3994 Advanced parallel and embedded computer systems project. This work was supported by The Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project IT4Innovations excellence in science - LQ1602 and by the IT4Innovations infrastructure which is supported from the Large Infrastructures for Research, Experimental Development and Innovations project IT4Innovations National Supercomputing Center - LM2015070.

VII. References

[1] Y-F. Zhou, A. Syed Arbab, and R.X. Xu. 2011. High intensity focused ultrasound in clinical tumor ablation. World journal of clinical oncology 2, 1 (2011), 8–27.

[2] G. Ter Haar, I. Rivens, L. Chen, and S. Riddler. 1991. High intensity focused ultrasound for the treatment of rat tumours. Physics in Medicine and Biology 36, 11 (1991), 1495–1501.

[3] M. Ichihara, K. Sasaki, S.I. Umemura, et al. 2007. Blood flow occlusion via ultrasound image-guided high-intensity focused ultrasound and its effect on tissue perfusion. Ultrasound in Medicine and Biology 33, 3 (2007), 452–459.

[4] V. Suomi, J. Jaros, B.E. Treeby, and R. Cleveland. 2017. Full Modelling of HighIntensity Focused Ultrasound and Thermal Heating in the Kidney using Realistic Patient Models. IEEE Transactions on Biomedical Engineering PP, 99 (2017), 1–12.

[5] N. Hansen and A. Ostermeier. 1996. Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In Proceedings of IEEE International Conference on Evolutionary Computation. IEEE, 312–317.

[6] S.A. Sapareto and W.C. Dewey. 1984. Thermal dose determination in cancer therapy. Int. Journal of Radiation Oncology, Biology, Physics 10, 6 (1984), 787–800.

[7] B.E. Treeby, J. Jaros, A.P. Rendell, and B.T. Cox. 2012. Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. The Journal of the Acoustical Society of America 131, 6 (2012), 4324–36.

[8] Cudova, Marta & E. Treeby, Bradley & Jaros, Jiri. (2018). Design of HIFU treatment plans using an evolutionary strategy. 1568-1575. 10.1145/3205651.3208268.

[9] H. Pennes. 1948. Analysis of tissue and arterial blood temperatures in the resting human forearm. Journal of applied physiology 1, 2 (1948), 93–122.