2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)
17-20 October, 2021. Melbourne, Australia

Distributed Evolutionary Design of High Intensity Focused Ultrasound Treatment Plans*

Jakub Chlebik[1] and Jiri Jaros[2]

Abstract— **High-Intensity Focused Ultrasound (HIFU) is a modern and still evolving technique used to treat a variety of solid malignant cells in a well-defined volume including breast, liver, pancreas, prostate or uterine fibroids or other general soft-tissue sarcomas. HIFU treatments allow a noninvasive and non-ionising approach when compared to more conventional cancer treatments, such as radio and chemo-therapy or open surgery, which can lead to a multitude of complications after the treatment. In recent years, a realistic thermal model accounting for patient-specific materials to design HIFU treatment plans was introduced, along with an evolutionary strategy to optimize them. However, the execution times of this model is too prohibitive to allow for a routine use. This paper presents a comparison of two distinct distributed evolutionary models employing a further optimized fitness model. The experiments show up to 6 times decrease in the evolution time. These improvements allowed to investigate a new real-life based benchmark and use-case.**

I. INTRODUCTION

The principle of the HIFU treatment is coagulative thermal necrosis by raising the temperature in the focal region by tens of degrees. The danger of potential tissue ablation outside of the focal region is low since the cytotoxic levels of temperature only occur in small volumes (about $1\,\mathrm{mm}$ by $10\,\mathrm{mm}$). Tumours of larger size are destroyed by chaining multiple focal regions, thus creating a contiguous lesion lattice encompassing the malignant tissue [1].

Compared to traditional methods such as radiotherapy, the method suffers from low delivery precision. Only with recent advancements in numerical methods, high performance computing, and physically relevant models, detailed simulations accurately capturing the relevant physical behaviour of HIFU waves and temperature distribution in a heterogeneous medium are now possible [2], [3].

This paper furthers the efforts towards the automated design of the treatment plans by taking another step in the fitness function optimization, and by use and comparasion of the already proposed *Island evolution* model with a simple *Farmer-Workers* model, using an equal size population. Experiments are then repeated on a newly introduced realistic use-case.

*This work was supported by The Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project "IT4Innovations National Supercomputing Center – LM2018140". This work was supported by Czech Science Foundation project 19-10137S.

[1]Jakub Chlebik and [2]Jiri Jaros are with the Centre of Excellence IT4Innovations, Faculty of Information Technology, Brno University of Technology, Brno, Czech Republic ichlebik@fit.vut.cz, jarosjir@fit.vut.cz

II. RELEVANT WORK

Recent work by Jaros [3] achieved first promising results at evolving precise HIFU treatment plans using *Covariance Matrix Adaptation Evolutionary Strategy* (CMA-ES) [4] with a physically correct fitness function. However, the process of evolution took up to 48 hours on very powerful computing machines. Additional effort on accelerating this computation was undertaken by Kuklis in [5], showing promising results comparing multitude of Island based evolution strategies to find the best parameter set-up. This paper presents a comparing review of the best found *Island* evolution strategy against the standard *Farmer-Workers* model, using a new, better optimized thermal model and a more demanding realistic HIFU surgery use case.

III. PROPOSED EVOLUTIONARY STRATEGY

This section describes the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), developed by N. Hansen [4]. After a brief introduction, the treatment plan encoding is outlined. Finally, the fitness function, based on the tissue-realistic heat diffusion model developed as part of the k-Wave toolbox is introduced [2], [6].

A. Covariance Matrix Adaptation Evolutionary Strategy

The Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) is a state-of-the-art real domain optimization method very popular for its ability to solve nonlinear, nonconvex objective functions. In contrast to the canonical Evolutionary Strategy, which moves the population towards the optimum by a random step size taken from a distribution with a fixed standard deviation, CMA-ES extends this approach by using a covariance matrix to describe the pairwise dependencies between genes [7].

In CMA-ES, a new population of λ new search points is generated by sampling a multivariate normal distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{C})$. This distribution is determined by its mean $\boldsymbol{\mu} \in \mathbb{R}^N$ and its symmetric and positive defined covariance matrix $\boldsymbol{C} \in \mathbb{R}^{N\times N}$. Together, these parameters create an ellipsoid on the search space, with the covariance giving it its shape and the mean describing its center, effectively creating a currently searched subspace. The length of the step is controlled by the step-size parameter $\sigma \in \mathbb{R}^N$. New generation is consequently sampled by the following formula:

$$\boldsymbol{x_i} \sim \boldsymbol{\mu} + \sigma\mathcal{N}(0, \boldsymbol{C}) \quad \text{for i} = 1, \ldots, \lambda \tag{1}$$

These newly generated individuals are first sorted by fitness and the best m of those individuals are selected.

Next, the mean value, step size and the covariance matrix are updated. The mean μ is updated by weighting intermediate recombination, where the weight of every selected individual is proportional to its rank in the population. The step size σ is controlled by the so-called evolution path. Conceptually, the evolution paths are the search paths the strategy takes over a number of generation steps. The main benefit of the CMA-ES is a relatively small population and fast convergence for real-valued problems. The step-size control allows for log-linear convergence and possibly linear scaling with dimension [7].

B. Encoding of the Treatment Plans

The treatment of large target areas using HIFU requires multiple sonications to cover the desired volume. The candidate solution I describes the trajectory the HIFU transducer follows. The treatment is not continuous but proceeds at precisely defined points in the tissue. The amount of energy delivered during a single sonication is given by the length of the sonication, represented by the variable t_{on}, and the length of the subsequent cooling interval, represented by the variable t_{off}. The number of sonications is limited to N, usually low to mid tens. If we consider 2D problems, one sonication can then be defined as a 4-tuple consisting of two spatial coordinates of the beam focus, and the temporal sonication and cooling intervals:

$$S_i = (x_i, y_i, t_{on}, t_{off}) \tag{2}$$

The entire solution is then defined as a sequence of sonications:

$$I = (S_1, S_2, \ldots, S_N) \tag{3}$$

With this considered, the optimization of the treatment plan is defined as the search for the optimal positions and sonication lengths for the given number of ablations to destroy the targeted volume of a malignant tissue while preserving the surrounding areas.

C. Fitness Function

With problem definition set, we need to describe the fitness function used to assess the quality of candidate plans. Such assessment is composed of several stages, described in the following sections [3].

1) ***Heat Deposition****:* The first step is to determine the heat energy deposition for every sonication using the predicted shape and position of the ultrasound focus and the sonication interval. This can be determined by a several complex numerical models, such as [6], [8], [9]. Unfortunately, their execution times are often too prohibitive for the use in evolutionary algorithms. Consequently, a few simplifications assuming a perfect transducer had to be made. (a) The center of the focus is possible to be placed at coordinates given by the sonication. (b) The distribution of the energy around the focus follows the Gaussian distribution [10]. For the sake of simplicity, we always assume the transducer axes to be aligned with the domain axes and only consider problems in 2D.

2) ***Thermal Model Execution****:* The second step is the execution of the numerical thermal model to calculate the temperature distribution in the domain during the treatment. The heat diffusion is modelled by the Pennes' bioheat equation [11] incorporating various tissue properties and effects of blood perfusion. Blood perfusion in particular is important to take into account since the blood vessels carry away considerable amounts of energy and thus cool down the tissue and prevent the ablation of malignant cells.

3) ***Thermal Damage Computation****:* Finally, the resulting tissue damage is evaluated using the *Sapereto-Dewey isoeffect thermal dose relationship* [12]. This relationship, called $CEM43$ [13], represents the equivalent time which would produce the same biological effects at a temperature of $43°C$. Generated spatial map of $CEM43$ is compared to a value of 240 (thermal dose of 240 minutes at $43°C$ irreversibly damages and coagulates critical cellular proteins, leading to tissue destruction) to produce a binary mask of destroyed tissue. Eq. (4) describes this process.

$$C_{xy} = \begin{cases} 0 & \text{if } CEM43_{xy} \leq 240 \\ 1 & \text{else} \end{cases} \tag{4}$$

To give the optimization algorithm more freedom, some "do not care" areas can be specified as well. Eq (5) describes the calculation of the resulting fitness value.

$$f = \int_0^Y \int_0^X ((D \cdot \overline{C}) + (P \cdot C))dxdy \tag{5}$$

X and Y are the domain sizes along the x and y axes, respectively. D is a target map specifying the area to be treated, C is a binary mask representing the actual result of the evaluated plan. $\overline{C}$ is a complementary mask for C, representing the nontreated area. P represents the prohibited area. As D and P are provided by the user and are specific to each patient, a professional user can specify the level of urgency for any given point by increasing its D or P values.

IV. PROPOSED IMPROVEMENTS

To further the efforts from [5], the entire optimization process was rewritten from Matlab to C++, eliminating the overhead of interfacing between two systems and allowing full control over the process, including the generation of heat energy matrices for deposition. Next, further optimizations of the source code were carried out. This concerns mostly better exploitation of a multi-core computation and vectorization using OpenMP implementation to allow for better scaling. Finally, the C++ code was compiled with the highest optimization level specific for the target CPUs.

A. Performance Analysis

With these changes, a new scaling test using the same hardware as [5] - 2 $\times$ Intel Xeon E5-2680v3 with 12 cores at 2.5 GHz - was performed, see Fig. 1. The results show the execution time decreasing rapidly up to 12 cores, from thereafter, a slight slow down can be observed, caused comunicatoin between both processors. However, this is a

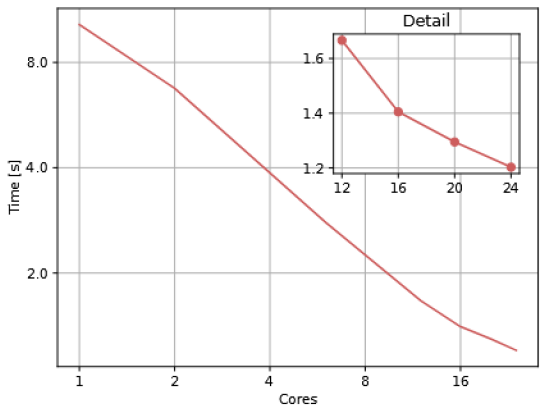

Fig. 1: Median of 15 execution times of a random single individual executed on 1 to 24 cores. One individual represents a treatment plan of 6 sonications. Log_2 scale on both axes.

major improvement in terms of scalability over the results presented in [5].

B. Distributed Evolution Models

The probability of finding the optimal solution by CMA-ES is known to increase with the population size [14], especially for non-linear problems. Faster evolution can be achieved by concurrent evaluation of multiple individuals. One option is to simply distribute the fitness evaluation to multiple computing nodes and gather results. This is called the *Farmer-Workers* model. An alternative is to use multiple smaller populations to run in parallel and exchange information periodically. This model is called the *Island* model of evolution. In this approach, each island employs several processor cores. These cores can be used either to evaluate multiple individuals concurrently, or to accelerate the evaluation of any single individual. The updated execution profile revealed only a 1.5% overhead introduced by the island model. Both models were implemented using the OpenMPI library.

*1) **Island Model**:* The driving idea behind the Island model is to use multiple smaller populations and increase diversity by accepting solution candidates evolved independently on other islands, and thus avoid local minima. Pseudocode displayed in Listing 1 shows the Island CMA-ES workflow. After the initialization, the evolution runs in a loop until the stopping condition is met. Each generation starts with evaluation of λ newly sampled individuals from the local CMA-ES, updating its internal state right after. However, every N generations migration occurs. During the migration, the population on each island is sorted according to its fitness and M best individuals are distributed to some other islands. The receiving islands accept these immigrants only if the acceptance condition is met. Next, P immigrants undergo a selection process to replace P worst individuals in the local population. At the end of each loop, the local CMA-ES models are updated. Some additional controls are employed to prevent too steep drift of the search space

```
1  def island():
2    while evo:
3      pop = sample_population(params, lambd)
4      fitVals = fitness_function(pop)
5      if migrateInterval(gen):
6        sortedI = index_sort(fitVals, >)
7        broadcast(fitVals[1:nBest])
8        (incPop, incFit) = recieve()
9        if accept(incFit):
10         for i in (1:nImmigrants):
11           incIndex = select(incFit)
12           orgIndex = sortedI[i]
13           pop[orgIndex] = incPop[incI]
14           fitVals[orgIndex] = incFit[incI]
15     end = cmaes_update(pop, fitVals)
16     gen += 1
17     evo = all_reduce(end, +)
```

Listing 1: Pseudocode of the Island CMA-ES.

```
1  def farmer():
2    while evo:
3      pop = sample_population(params, lambd)
4      popPart = scatter(pop)
5      fitVals = fitness_function(popPart)
6      append(fitVals, gather_results())
7      evo = cmaes_update(pop, fitVals)
8      gen += 1
9      broadcast(evo)
10
11 def worker():
12   while evo:
13     popPart = gather()
14     fitVals = fitness_function(popPart)
15     send(fitVals, ROOT)
16     evo = recieve()
```

Listing 2: Pseudocode of the Farmer-Workers CMA-ES.

ellipsoid, limiting the ability of CMA-ES to hard pivot the evolution paths based on a single individual outside of its distribution. Before the next loop repeats, each island checks whether the optimization met the stopping condition and distributes this information accordingly.

*2) **Farmer-Workers model**:* Two complementary pseudocodes of this approach are displayed in Listing 2. In this model, the farmer node (see Listing 2) runs the optimization process, generating a large population of individuals. The farmer then distributes the computation of fitness evaluations among workers. After a worker is done with the assigned job batch, the results are sent back to the farmer node. Once all workers complete their tasks, the farmer updates the internal state of the optimization algorithm and repeats the process with next sample of individuals. Once the stopping criterion is satisfied, the farmer node informs workers and terminates. The advantage of this approach is a larger and thus more diverse population.

V. EXPERIMENTS

This section details the setup for our experiments. The treatment setup parameters and the parameters for the distributed evolution models - the Farmer-Worker (*FW*) model and the Island (*I*) evolution model. For completeness, the ordinary one node sequential (*SQ*) model is also compared. 15 independent runs for each model were executed on two

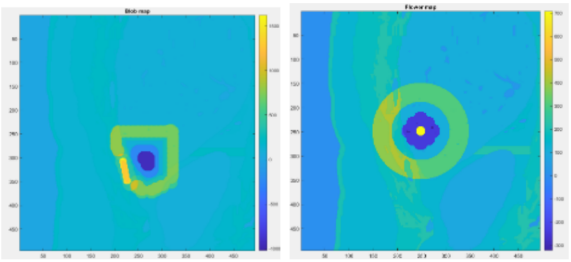

Fig. 2: A visualisation of the target and penalization masks used in benchmarks. Medium data are taken from the open-source AustinWoman voxel model. Left picture represents the monolithic use case. Right pictures is the flower use case.

different benchmarks. To compare the two presented models of distributed evolution, two metrics were examined:

- The success rate - The percentage of all runs capable of finding the optimal solution.
- The time to solution - The computational model capable of delivering satisfying results faster should be favoured for further development.

A. Treatment Setup

To set as realistic conditions as possible, the same benchmark as presented in the paper [3] was used. It represents the use case of ablation of one abdominal target within the right lobe of the liver. Tissue parameters were acquired from the open-source AustinWoman voxel model [15]. Multiple levels of severity for both the target and the penalization maps were employed to diversify the search space. Two levels to represent the target area D and three levels for the prohibited area P.

Additionally, thanks to the improvements in performance of fitness evaluation, we present a new benchmarking case - a tumour overgrown around a critical, healthy tissue, such as arteries, large veins or even urethra. For this case study, a very rigid penalization mask was used, with the aim of preventing ablation of the area in the middle of the target at all costs. A wider "do not care" area around the target was also employed to allow for more leeway in finding the outer shape. This use case model was named a 'flower'.

Both benchmarks, can be seen in Fig. 2.

As for the heat source; the spatial peak of the volume rate of heat deposition was set to $100\,\mathrm{W\,cm^{-3}}$, which approximately matches the clinical data.

The parameters of the numerical thermal model were set according to numerical convergence testing as follows:

- Domain size 495×495 grid points, periodic boundary conditions, with $0.2\,\mathrm{mm}$ spatial resolution.
- Temporal resolution of heat deposition calculations was set to $0.1\,\mathrm{s}$.
- Ultrasound focus center position limited to the bounding box $[270, 230] \times [345, 295]$ for the monolithic case study and $[287, 207] \times [287, 207]$ for the flower case study.
- Temporal intervals are bound to $[0,\ 20\,\mathrm{s}]$ for both the monolithic and the flower case. However, the flower model used a wider - $[0,\ 150\,\mathrm{s}]$ - interval for the cooling axis. This change is introduced because of the heat diffusion into the surrounding tissue after a sonication, causing energy accumulation inside the sonications area overlap. The aim of this change is to introduce a possible cooling of the surrouding non-targeted tissue.
- 6 ($6\times$ in figures) and 20 ($20\times$) sonications used for both benchmarks. Setting the number of sonications correcly is a very tedious task requiring expert knowledge. The success rate and evolution time are highly dependant on the number of sonications used. In the paper [3], treatment plans with varying numbers of sonications have been investigated, with the conclusion of 6 and 8 being the best. We decided to use 6 to be able to compare with [5] and then introduce a benchmark using 20, to see the behaviour of the treatment planning methods on a higher sonication use case, expanding the original study with a real world scenario use case. This test is relevant as a proof of concept as it was not previously possible.

B. Optimization Run Parameters

One of the big advantages of CMA-ES is the lack of parameters necessary to be tuned for a successful evolution. This advantage is preserved using the Farmer-Workers model, however, the proposed Island model introduces several parameters heavily influencing the behaviour of the evolution. We are building on a paper [5] which already did a deep study on the effect of various Island evolution parameters on the ability to solve this specific problem. As such, we adopted the results and aim to expand the research.

- Number of Islands - The island model with 6 nodes (x6) was investigated. Since the performance improvements allowed to scale past 12 cores and the reference machine uses 2×12 cores per node, we decided to fully levarege the fitness scaling and use all 24 cores for one island. Thus, 6 islands per run were used.
- Size of population per island - The size of the population is 13 and 17 individuals for 6 and 20 sonications, respectively. These are default values decided by CMA-ES.
- Migration interval - The decision to migrate every $M = 10$ generations was made after brief evaluation. This seems like a reasonable compromise [16] as too frequent migration causes a loss in diversity and too rare slows down convergence.
- Island topology - We decided to adopt a fully connected graph topology from [5]. Based on our empirical data, there was no or negligible difference with a ring or even randomly connected topology.
- Acceptance policy - Each island has a probability to accept new individuals into the population based on the current state of optimization. We decided to adopt the acceptance policy introduced in paper [5] as it showed a good improvement in convergence times.
- Replacement strategy - Each island receives more immigrants than it can reasonably keep in its population.

Which migrants are integrated is decided by a tournament selection.

The population size for the *FW* model is adjusted to be the same as the sum of the population sizes across all islands in the competing model. This means 78 individual for 6 sonications (6×13) and 102 (6×17) for 20 sonications. The population size for the *SQ* model is the default population determined by CMA-ES and serves as a comparison of distributed models to a standard single node search process. Lastly, the evolution algorithm is allowed to run for maximum of 8 hours to simulate the result urgency requirement.

VI. RESULTS

A. Monolithic Target

Different distributed evolution models were evaluated using the introduced monolithic case study. Treatment plans consisting of 6 and 20 sonications were measured and the final fitness of the solutions plotted in Fig. 3. The upper plot describes the results of planning for the monolithic case, the lower represents the flower case.

We can see that for lower numbers of sonications, the island evolution model provides superior results, as all of its runs managed to find the optimal solution, confirming results from [5]. However, from the same figure we can observe that the Farmer-Workers model is not very far behind.

Since the median of the *FW* model settles to the optimal fitness value, more than 50% of all runs managed to find an optimal solution. Furthermore, when analysing Fig. 4, we can see that when it did, it generally converged faster than the Island model. It should also be mentioned that very small fitness values (less than 50) are usually acceptable solutions too, as only a small fraction (less than 0.5%) of the targeted volume was not treated. Alternatively, a very small volume of penalized protected area was destroyed.

For the higher number of sonications, the FW model pulls ahead of the I model substantially, boasting 100% success rate with a much faster time to solution, compared to the Island model, which managed to find the optimal solution in less than 25% of all runs.

B. Flower

Solving the flower benchmark proved to be more difficult than the original monolithic one. None of the runs managed to find the optimal treatment plan for both 6 and 20 sonications. Nevertheless, smaller numbers of sonications appear to be better suited for solving this benchmark. If the assumptions are correct, the reason for this is that fewer sonications provide less residual heat to diffuse into the prohibited middle part. However, a smaller dimension of the problem, with all other parameters remaining the same, allows for faster heat diffusion simulations and thus a higher number of fitness evaluations inside the given time window. To confirm that the number of sonications is the culprit of this behaviour, we plot the fitness against the number of fitness evaluations for both options (Fig. 5). In this figure, we can see that better results are not because of the higher number of fitness evaluations. Worthy of mention might be the fact that the median fitness of around 70 for the FW6x6 model might be acceptable for treatment in some edge cases. Penalization mask used for this benchmarks causes even a single pixel hit of the middle area to highly increase fitness.

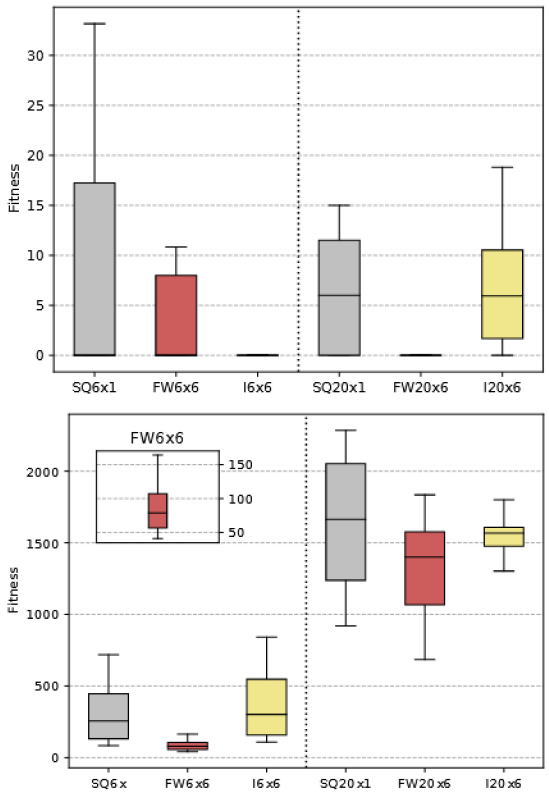

Fig. 3: The final fitness value of the treatment planning process for 6 and 20 sonications on the monolithic target (top) and flower target (bottom).

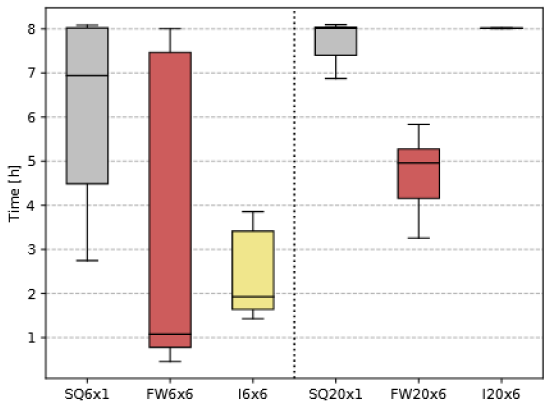

Fig. 4: The time needed for each model to find the solution for the monolithic target. All runs were capped at 8 hours. Times for the flower benchmarks are not shown because of the inability to find the optimal value and thus always running for the maximum allowed time.

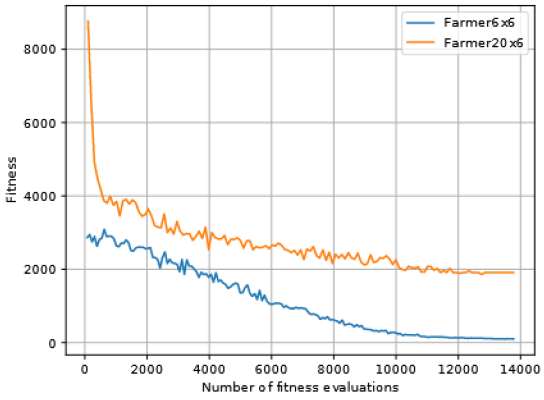

Fig. 5: Median fitness of best solutions at given number of fitness evaluations for both number of sonications. The data for option 6x are cut and aligned to the final number of fitness evaluations for 20x.

VII. CONCLUSIONS

This paper has presented the results of our latest work on optimizing HIFU treatment plans. Farmer-Workers model seems to, in general, outperform the Island based approach, most likely thanks to exposing the entire population data to the CMA evolution strategy, instead of just the best individuals from separately evolved populations. It appears that a simple increase in population diversity by injection of separate individual is not enough. Considering the approach CMA-ES uses to improve solutions, a possible reason for this could be that randomly perturbing the population from outside of the distribution slows down the exploitative cabalities of the method, forcing an adaptation of the entire covariance to a solution that might be from a different local optima or on the edge of the search space. The latter is known from literature to harm the performance of CMA-ES [14]. Meanwhile, for large populations governed by a single instance of the algorithm, the method has access to a much larger sample size to learn all pairwise dependencies correctly. However, thanks to a newly optimized model, we were able to move a previous computational time limit of 48 hours down to 8 hours while still finding feasible solutions without any loss in precision. Furthermore, introducing the distributed computation models allows us to find a good enough solution for the basic monolithic use case in less than 2 hours. This in turn enabled the use of another realistic setting with a higher number of sonications - 20.

Increasing the amount of sonication (hundreds or more), however, raises the dimension of the solved problem. Although our data shows that 20 sonications are still solvable in roughly 5 hours, we must consider the inevitable transition to 3D spatial domain. This will further increase the complexity of finding a feasible solution, by at the very least introducing one more variable per sonication. While we can compensate by also increasing the number of computational nodes used, this approach quickly becomes cost inefficient and unsustainable. Most probable solution here lies in a change of our heat diffusion modeling, e. g. a GPU acceleration or a completely new approach, based on recent advancements in machine learning or surrogate modeling.

Sadly, the introduced flower target still requires more work to be done before moving further to 3D. Despite the fact that the solution for 6 sonications shows promise, the 20 sonication benchmark still provides very unsatisfactory results. The most likely culprit here is the heat accumulation inside the protected region. One possible solution to this would be an introduction of yet another variable into the search space - the transducer power. Another could be automaticaly adapting the number of sonications during the optimization. Both of these, however increase the complexity of the entire optimization process, even more feeding into dimension creep problem outlined above.

REFERENCES

[1] M. Ichihara, K. Sasaki, *et al.*, "Blood flow occlusion via ultrasound image-guided high-intensity focused ultrasound and its effect on tissue perfusion," *Ultrasound in Medicine and Biology*, vol. 33, no. 3, pp. 452–459, Mar. 2007.

[2] V. Suomi, J. Jaros, *et al.*, "Full modeling of high-intensity focused ultrasound and thermal heating in the kidney using realistic patient models," *IEEE Transactions on Biomedical Engineering*, vol. PP, pp. 1–1, 09 2018.

[3] M. Cudova *et al.*, "Design of hifu treatment plans using an evolutionary strategy," ser. GECCO '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1568–1575. [Online]. Available: https://doi.org/10.1145/3205651.3208268

[4] N. Hansen, *The CMA Evolution Strategy: A Comparing Review*, 06 2007, vol. 192, pp. 75–102.

[5] F. Kuklis, M. Jaros, and J. Jaros, "Accelerated design of hifu treatment plans using island-based evolutionary strategy," *Applications of Evolutionary Computation Lecture Notes in Computer Science*, p. 463–478, Apr 2020.

[6] B. Treeby, J. Jaros, *et al.*, "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*, vol. 131, pp. 4324–36, 06 2012.

[7] N. Hansen, "The cma evolution strategy: a tutorial," 01 2010.

[8] J. Huijssen and M. D. Verweij, "An iterative method for the computation of nonlinear, wide-angle, pulsed acoustic fields of medical diagnostic transducers," *J Acoust Soc Am*, vol. 127, no. 1, pp. 33–44, Jan 2010.

[9] F. Marquet, M. Pernot, *et al.*, "Non-invasive transcranial ultrasound therapy based on a 3D CT scan: protocol validation and in vitro results," *Phys Med Biol*, vol. 54, no. 9, pp. 2597–2613, May 2009.

[10] J. Jenne *et al.*, "High-intensity focused ultrasound: Principles, therapy guidance, simulations and applications." *Zeitschrift fur medizinische Physik*, vol. 22, 08 2012.

[11] H. H. PENNES, "Analysis of tissue and arterial blood temperatures in the resting human forearm," *J Appl Physiol*, vol. 1, no. 2, pp. 93–122, Aug 1948.

[12] S. A. Sapareto and W. C. Dewey, "Thermal dose determination in cancer therapy," *Int J Radiat Oncol Biol Phys*, vol. 10, no. 6, pp. 787–800, Jun 1984.

[13] Y. Zhou, "High intensity focused ultrasound in clinical tumor ablation," *World journal of clinical oncology*, vol. 2, pp. 8–27, 01 2011.

[14] N. Hansen and S. Kern, "Evaluating the cma evolution strategy on multimodal test functions," in *Parallel Problem Solving from Nature - PPSN VIII*, X. Yao, E. K. Burke, J. A. Lozano, J. Smith, J. J. Merelo-Guervos, J. A. Bullinaria, J. E. Rowe, P. Tino, A. Kaban, and H.-P. Schwefel, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 282–291.

[15] J. Massey, C. Geyik, *et al.*, "Austinman and austinwoman: High fidelity, reproducible, and open-source electromagnetic voxel models," 06 2012.

[16] Z. Skolicki and K. De Jong, "The influence of migration intervals on island models," 01 2005, pp. 1295–1302.