



Accelerated Design of HIFU Treatment Plans Using Island-Based Evolutionary Strategy

Filip Kuklis, Marta Jaros , and Jiri Jaros  

Faculty of Information Technology, Centre of Excellence IT4Innovations,
Brno University of Technology, Bozetechova 2, 612 66 Brno, Czech Republic
jarosjir@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. Such a treatment consists of a series of short sonications destroying small volumes of tissue. High-quality treatment plans allow to precisely target malignant tissue and protect surrounding healthy tissue. Recently, we developed an evolutionary strategy to design such HIFU treatment plans using patient-specific material properties and a realistic thermal model. Unfortunately, the execution time was prohibitive for routine use. Here, we present an optimized evolutionary strategy based on island model parallelization and a revised fitness function implementation. The proposed improvements allow to develop a good treatment plan 4 times faster and with 5% higher success rate.

Keywords: Evolutionary strategy · Island model · HIFU · Treatment planning · k-Wave toolbox

1 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 [24].

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 1 mm in diameter and about 10 mm in length), which minimizes the potential for thermal damage to tissue outside the

focal region. 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. Blood perfusion may also carry away a significant amount of energy and deteriorate the treatment outcome [12].

Despite the advantages of HIFU, this method still suffers from delivery precision in contrast to other established therapies such as radiotherapy. 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 [21]. However, model-based treatment planning (determining the best transducer position and sonication parameters to deliver the ultrasound energy to the planning target volume) is still currently performed in a relatively rudimentary way based on heuristics rather than physical models of the therapy. This leads to rather poor quality of the treatment plans.

Recently, first steps towards the automated design of precise HIFU treatment plans have been made via the use of Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) [8] in combination with a physically relevant fitness function [4]. The planning algorithm achieved promising results by producing treatment plans with negligible mistreated and undertreated areas. Unfortunately, the evolution process often took more than one day, even on very powerful computing servers integrating two processors with 24 cores in total.

This paper presents our effort in reducing the evolution run time. It was decided to adopt the island model of EA [1, 2, 20] where the population is split into several sub-populations assigned to particular computational resources. These sub-populations mostly evolve independently, which ensures higher diversity of the evolutionary process, yet share some collective knowledge about promising areas in the search space, which improves the convergence towards global optima. Furthermore, since the fitness function evaluation is extremely time-consuming, considerable time has been spent in the code optimization.

The rest of the paper is organized as follows. The next section recapitulates the structure of the evolutionary algorithm, its encoding and fitness function. Section 3 details the optimization of the fitness function. Section 4 describes the island model implementation. The parameters of the island model and the quality of proposed optimizations are elaborated in Sect. 5. The last section concludes the paper and draws future work directions.

2 Proposed Algorithm

This section first describes the optimization algorithm based on the Matlab implementation of the Covariance Matrix Adaptation (CMA) Evolutionary Strategy (ES) developed by Hansen [7]. Then, the treatment plan encoding is outlined. Finally, the fitness function based on the tissue-realistic heat diffusion developed as part of the k-Wave toolbox is introduced [21, 22].

2.1 Evolutionary Strategy

The CMA-ES [7, 8] is a very popular stochastic method for real-parameter (continuous domain) optimization of nonlinear, non-convex objective functions. The CMA [9] describes the pairwise dependencies between variables/genes on the top of the classic ES.

In CMA-ES, a population of λ new search points (individuals, offspring) is generated by sampling a multivariate normal distribution $\mathcal{N}(\mathbf{m}, \mathbf{C})$ determined by its mean $\mathbf{m} \in \mathbb{R}^N$ and its symmetric and positive defined covariance matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$, which determines the shape of the distribution ellipsoid. The length of the step is controlled by the so-called step-size parameter $\sigma \in \mathbb{R}^N$:

$$x_i \sim \mathbf{m} + \sigma \mathcal{N}_i(0, \mathbf{C}) \text{ for } i = 1, \dots, \lambda. \quad (1)$$

The newly generated individuals are first ranked according to their fitness and then the μ individuals are selected. The elitism is not used. Next, the mean value, step size and the covariance matrix are updated. The mean value \mathbf{m} is updated by weighted intermediate recombination where the weight of every selected individual is proportional to its rank. The CMA-ES utilizes an evolution path to control the step size σ . Conceptually, the evolution path is the search path the strategy takes over a number of generation steps. The adaptation of the covariance matrix follows a natural gradient approximation of the expected fitness. The adaptation procedure first learns all pairwise dependencies between all variables. Then, it conducts a principal components (eigenvectors) analysis (PCA) of steps sequentially in time and space. Finally, a new rotated problem representation is determined using the Mahalanobis metric [3].

The main benefit of the CMA-ES is a very small population and fast convergence for real-valued problems compared to Genetic Algorithms (GA) [2] or Estimation of Distribution Algorithms (EDA) [14]. The step-size control facilitates fast (log-linear) convergence and possibly linear scaling with the dimension. The covariance matrix adaptation increases the likelihood of previously successful steps and can improve performance by orders of magnitude [8].

2.2 Treatment Plan Encoding

The ablation of large target areas using HIFU requires multiple sonications to effectively cover this area. The candidate solution I describes the trajectory the HIFU transducer follows in the tissue during the treatment. The treatment is not continuous but proceeds at precisely defined points in the tissue where the HIFU focus is placed. The number of sonications is limited to N , usually low tens. The amount of energy delivered during a single sonication is given by the length of the sonication t_{on} and the length of the subsequent cooling interval t_{off} . One sonication can thus be defined as a 4-tuple S_i composed of two spatial coordinates of the beam focus (only 2D problems are considered), and sonication and cooling intervals t_{on} and t_{off} , respectively:

$$I = (S_1, S_2, \dots, S_N), \text{ where } S_i = (x_i, y_i, t_{\text{on}, i}, t_{\text{off}, i}) \quad (2)$$

2.3 Fitness Function

Generally, the treatment planning problem is defined as the search for suitable positions and sonication times for the specified number sonications to cover the target area and minimize the volume of mistreated areas. The assessment of the treatment plan quality is composed of several stages, detailed in [4]. First, the heat deposition for every sonication is determined using the predicted shape and position of the ultrasound focus and the sonication length. Physically, they can be determined by complex numerical models [11, 15, 22]. However, since their execution times are often prohibitive for applications in evolutionary algorithms, several simplifications had to be made: (1) the centre of the focus can be placed at coordinates given by the sonication $S_i = [x_i, y_i]$, and (2) the distribution of the energy in the focus follows the Gaussian distribution [5, 23].

Second, the numerical thermal model is executed on the whole sequence of sonications to calculate the temperature distribution in the domain during the treatment. The heat diffusion is modelled by the Pennes' bioheat equation [18] which has a corresponding ultrasound energy absorption term as a source term, incorporates various tissue properties and effects of blood perfusion.

Figure 1 shows the spatial temperature distribution along the main focus axis at the end of the first sonication, and then every 20 s during the cooling period. Without the proximity of a large blood vessel, the spatial temperature distribution follows the Gaussian distribution. The temperature magnitude declines from a peak of 72 °C at the end of the sonication down to about 46 °C at the end of the cooling interval. On the other hand, the area with temperature exceeding 37 °C is slowly growing.

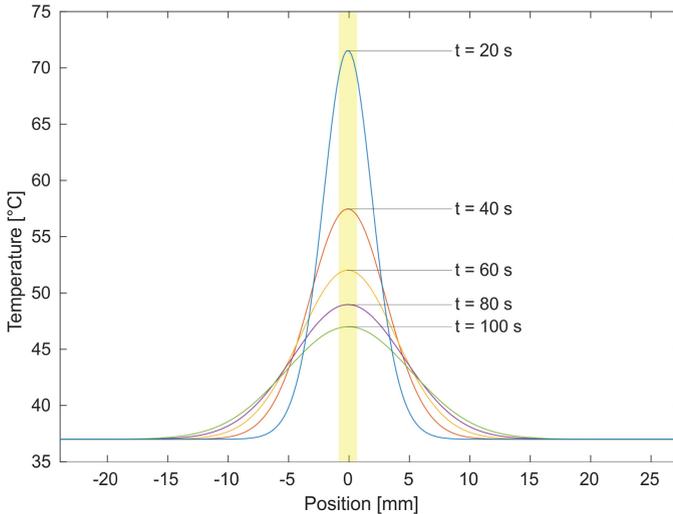


Fig. 1. Heat distribution along the main focus axis during a sonication with $t_{\text{on}} = 20$ s and $t_{\text{off}} = 100$ s at time $t = 20$ s, 40 s, 60 s, 80 s, 100 s. (Color figure online)

Thermal damage is computed using the Sapareto-Dewey iso-effect thermal dose relationship [19] which is expressed in seconds and represents the equivalent time which would produce the same biological effects at a temperature of 43 °C. This metric is called cumulative equivalent minutes at 43 °C (CEM_{43}). CEM_{43} is calculated for every point in the tissue and summed up over all sonications. Thermal doses of 240 min at 43 °C irreversibly damage and coagulate critical cellular protein, tissue structural components and the vasculature leading to immediate tissue destruction [24]. The area with the dose exceeding 240 CEM_{43} is depicted in Fig. 1 by a yellow bar.

The output of the thermal model is a spatial map of CEM_{43} accumulated over the whole treatment. This map is thresholded by a value of 240 to produce a binary mask of destroyed tissue. The evaluation of the quality of the HIFU treatment is based on the assumption that all tissue in the target area is destroyed while all tissue in the prohibited area (organs at risk) is left unharmed. In order to give the optimization algorithm some freedom, do not care areas can be specified as well.

The fitness function for a 2D case can be written as

$$f = \int_0^Y \int_0^X ((D * \bar{C}) + (P * C)) dx dy \quad (3)$$

$$C = \begin{cases} 0 & \text{for } CEM_{43} \leq 240 \\ 1 & \text{for } CEM_{43} > 240 \end{cases}$$

$$D \in \mathbb{R}^+, P \in \mathbb{R}^+,$$

where X, Y are domain sizes along the x and y axes, respectively, C is a binary mask representing the actually treated area, \bar{C} is a complementary mask representing the non-treated area, D is a target map specifying the area to be treated and P represents prohibited area. Since D and P are defined as functions over 2D space, users can specify the level of urgency a given point in the space shall be treated or protected with. The goal is then to minimize the fitness function.

3 Acceleration of Fitness Function

The probability of finding the optimal solution by CMA-ES is known to increase with the population size [10]. However, a larger population usually implies a much higher number of evaluations. As can be seen in Table 1, the fitness function is very complex due to the realistic simulation of heat diffusion which consumes over 99% of the computational time. That only allows populations with at most 40 individuals to be used due to a time constraint of 48 h. Therefore, the first step towards a more robust EA is to analyze and optimize the fitness function.

3.1 Analysis of the Matlab Implementation

The fitness function consists of the calculation of the heat diffusion followed by the assessment of the treated area. We only focused on the heat diffusion since the treated area assessment is computationally trivial.

Table 1. Time profile of the whole evolution process executed by Matlab.

Code	Calls	Total time	% Time
Fitness evaluation	2,990	47,998 s	99.9%
Population initialization	1	13.87 s	0.1%
Final solution evaluation	1	12.53 s	0.0%
Writes to log file	140	1.21 s	0.0%
Other		3.34 s	0.0%
Total		48,029 s	100%

The heat diffusion code was originally implemented in Matlab using the k-Wave toolbox [22]. Its pseudocode can be seen in Listing 1.1. This code supports precise tissue parameter settings derived from patient-specific models of the tissue anatomy discretized into a grid with spatial and temporal resolutions set according to the convergence testing, see Sect. 5.

The computation is based on a k-space pseudospectral scheme in which spatial gradients are calculated using the Fourier collocation spectral method and temporal gradients are calculated using a k-space corrected finite difference scheme. The precomputations of the k-space term and the heat source term are based on simple matrix operations and two Fourier transforms. Since being executed only once, this part of computation is negligible. On the other hand, the functions inside the loop are called between 50 and 200 times for typical sonication lengths. The computation of the divergence term uses 3 forward and 4 inverse 2D Fourier transforms. Hence, 350–1400 FFTs are executed. This causes the computation of divergence term to take almost 85% of the computation time of the whole simulation. The update of the damage integral takes about 10% and the rest of the operations amount to last 5%.

Listing 1.1. Pseudocode of the heat deposition calculation for one sonication.

```

Precompute k-space derivative term;
Precompute heat source term;

Loop through the sonication time span
begin
  Compute perfusion term;
  Compute divergence term;
  Update tissue temperature;
  Update damage integral;
end

```

3.2 Optimized Implementation

In order to maximize computational efficiency, the fitness function was rewritten from Matlab to C++ with several low-level optimizations. First, the Matlab FFT [6] taking about 75% of the execution time was replaced with the Intel MKL [13] version which offers 89% faster execution on the domain sizes of interest. Second, the element-wise matrix-matrix operations were parallelized and vectorized using

the C++ OpenMP 4.5 library [17] to exploit multiple processor cores and vector instructions such as Intel AVX. Multiple mathematical operations were applied to each grid point where possible to maximize temporal data locality. Finally, the C++ code was compiled with the highest optimization level tuned for the CPUs being used in our experiments.

Since the medical data processing and the CMA-ES itself was implemented in Matlab, the optimized fitness function was wrapped by a MEX function to be directly invoked from within Matlab without any additional overhead.

3.3 Performance Evaluation

The optimization of the fitness function itself brought more than twofold reduction of the execution time. In order to use an appropriate level of parallelism a scaling test was performed, see Fig. 2. The question to be answered was whether it is better to evaluate multiple individuals at the same time, or use multiple processor cores to evaluate a single one. Since CMA-ES uses a very small population and every individual may take a considerably different time, an even usage of all cores may become an issue. In the scaling test, the number of cores collaborating to evaluate a single individual was progressively increased from 1 to 24. The execution time of a single individual decreases almost linearly up to 12 cores (a single processor). When both processors are used, the performance deteriorates significantly. This can be attributed to the scaling of the Intel FFTs and the inter-CPU interconnect.

This measurement opens three almost equal possibilities how the population evaluation can be parallelized while keeping a reasonable level of load balance:

1. 6 individuals at the same time, each using 4 cores,
2. 3 individuals at the same time, each using 8 cores,
3. 2 individuals at the same time, each using 12 cores.

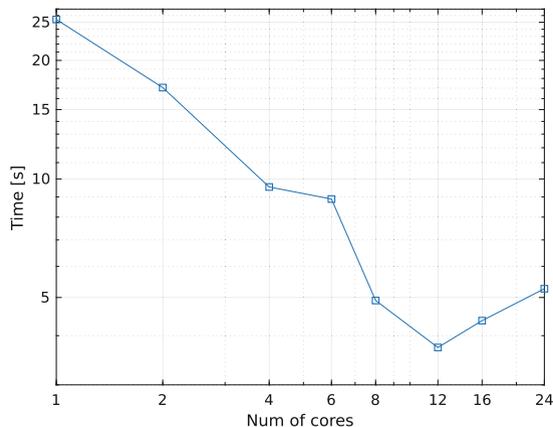


Fig. 2. Execution time of the fitness function on a dual-socket server with 24 cores.

4 Island Model

Faster evolution can be achieved by concurrent evaluation of multiple individuals. One possibility is to keep the evolutionary algorithm unchanged and only parallelize the evaluation. We, however, decided to split the population into multiple ones and run a parallel island-based CMA-ES with a subset of cores assigned to each island. These cores can be used either to evaluate multiple individuals concurrently, or to accelerate the evaluation of a single individual. The updated execution profile revealed only a 1.5% overhead introduced by the island model.

The pseudocode displayed in Listing 1.2 shows the island CMA-ES workflow. First, the algorithm is initialized. Then the evolution runs in a loop until the stop condition is met. The loop starts with generation and evaluation of λ new individuals from the local island CMA-ES model. The islands migrate individuals every M -th generation. During the migration, the population on each island is sorted according to the fitness values and N best individuals are broadcast to other islands. The other islands accept these individuals only if the acceptance condition is satisfied. Next, some of the immigrants are selected by a roulette wheel to replace N worst individuals in the local population. At the end of the loop, the local CMA-ES models are updated based on the newly formed population and the loop repeats.

Listing 1.2. Pseudocode of the island-based CMA-ES.

```

1  init_params(params);
2  while(stopflag)
3      % Generate and evaluate local population
4      for k = 1:lambda
5          ind = create_new_individual(params);
6          fitness(k) = eval_thermal_model(ind);
7      end
8
9      % Perform migration
10     if(mod(gen, M) == 0)
11         sort(fitness);
12         for i = 1:N % Islands perform migration in turn
13             if(i == current_island) % Who broadcasts individuals
14                 broadcast(i, fitness(1:n_best));
15             else if (accept_cond) % Shall I accept migrants
16                 rcv(i) = broadcast(i);
17             end
18             if(glob_select) % Global or per island selection
19                 sel(:) = roulette(rcv(:));
20             else if(isl_select)
21                 for i = 1:size(rcv)
22                     sel(i) = roulette(rcv(i));
23                 end
24             end
25         end
26         % Replace worst part of population
27         fitness(lambda - n_sel:lambda) = sel(:);
28     end
29
30     % Update local model
31     select_parents;
32     update_params(params, fitness);
33     gen = gen + 1;
34 end

```

5 Experimental Results

The experimental work presented in this paper investigates the benefits of several different versions of the island based CMA-ES and compares them with a pan-population (PP) version. To allow statistical evaluation, 15 independent runs for each version were executed with the maximum execution time per run limited to 48 h. To compare the original CMA-ES treatment planning with the island-based CMA-ES, two main metrics were examined: (1) success rate, (2) total number of fitness evaluations to converge. Success rate represents the percentage of runs which found the optimal solution with the fitness value of zero. In such a case, the treated area covers the whole desired area and the prohibited area is unharmed. The total number of fitness evaluations to converge is straightforward for the pan-population model. In the case of the island model, the total number of evaluation is the sum over all islands.

During the evaluation of the proposed algorithm we strove to work under as realistic conditions as possible. The same benchmark as in the paper by Cudova [4] was used. The HIFU treatment plans were designed for a representative map of the biological materials acquired from the open-source AustinWoman voxel model [16]. As a case study, one abdominal target within the right lobe of the liver was used. Two levels of D were chosen, a higher one ($D_{x,y} = 2$) in the middle of the target area, and a lower one ($D_{x,y} = 1$) close to the boundaries. The prohibited area P was marked by three different levels of importance. The highest one ($P_{x,y} = 5$) covered the rib and the tendon. The middle one ($P_{x,y} = 2$) covered the fat layers and the areas further from the treated areas. The lowest level ($P_{x,y} = 1$) was used for areas neighboring the treated areas. To make the interface between treated and prohibited areas smooth, a thin don't care area was used ($D_{x,y} = 0$ and $P_{x,y} = 0$). The size of the heat source was based on a single element transducer using the nominal properties of the HAIFU JC-200. The spatial peak of the volume rate of heat deposition was set to 100 W/cm^2 , which approximately matches the values used for clinical treatments.

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

- Discretized simulation domain size 495×495 grid points, periodic boundary condition.
- Spatial resolution 0.2 mm.
- Temporal resolution 0.1 s.
- The total length of the simulation $\sum_{i=0}^N (t_{\text{on},i} + t_{\text{off},i})$.
- Allowed positions of the ultrasound focus center limited to the bounding box at grid positions $[270, 230] \times [345, 295]$.
- Maximum sonication and cooling periods $t_{\text{on}} = [0, 20 \text{ s}]$, $t_{\text{off}} = [0, 20 \text{ s}]$.
- Number of sonications considered $N = 6$.

The success rate and evolution time is highly dependent on the number of sonications used in the treatment. Generally, the more the sonications is used, the easier the job to design an optimal treatment plan is. Naturally, it is easier to cover a given area with a higher number of smaller dots. On the other hand,

the applicaiton of the treatment takes much longer as well as the fitness function evaluation. This is given by a rising number of thermal model invocation in the simulation, and heating od distant points in the real treatment. In our previous paper [4], treatment plans with 4, 5, 6, 8 and 10 sonications have been deeply investigated. Six and eight sonications were concluded to be the best. Since shorter treatments poses a harder problem, we decided to use 6 sonications in the rest of the paper.

5.1 Examined Parameters of Island Model

Apart from CMA-ES related parameters, the proposed island model introduces several parameters influencing the behavior of the evolution. The following list elaborates the parameters of interest along with their examined values.

1. *Number of islands* - Having 24 computers cores and considering the scalability of the fitness function, we investigate the island CMA-ES with 3 and 6 islands and compared them with the original pan-population version.
2. *Total number of individuals* - The number of individuals per island is $\lambda_{\text{island}} = 13$, which is a default population size for CMA-ES with 6 sonications. The total population size for 3 and 6 islands reaches 39 and 78, respectively.
3. *Migration interval* - The number of generations between two migrations M remains constant during the evolution and is set to 1, 2 or 3 generations.
4. *Migration selection strategy* - Depending on the strategy, one or three best individuals are migrated.
5. *Island topology* - The island topology is a fully connected graph.
6. *Acceptance policy* - Each island has a predefined probability determining whether the immigrants are accepted or refused. This probability is progressively decreased along with the increasing quality of local individuals.
7. *Replace strategy* - Each island replaces R worst individuals in its population with R migrants. The migrants are selected using a roulette wheel.

5.2 Pan-Population Model

The results of the original PP model with the Matlab fitness function evaluation for with the population size $\lambda = 13, 20, 40$ were taken from the previous paper [4]. Since the optimized fitness function enables to shorten the computation time by a factor of two, another case with a population of $\lambda = 78$ individuals is added into the comparison.

5.3 CMA-ES with 6 Islands

In this section, four examined variants of the 6-island model with the total population of 78 individuals are explained.

1. *Send 1 Best - Accept All (6ISB)* - Every island broadcasts its best individual to the others. The others accept all migrants and replace the worse part of local populations.

2. *Send 3 Best - Select Roulette (6IS3)* - Three best individuals are broadcast but only a single immigrant from each island selected by a roulette wheel is accepted.
3. *Send 1 or 3 Best - Select Roulette - Acceptance Policy (6IS1AP or 6IS3AP)* - The strategy tries to maximize the diversity among islands by introducing the acceptance probability that decreases with the quality of local solutions, see Table 2. This allows the islands to converge to different optimal solutions which then can be displayed to the clinicians for deeper evaluation.

Table 2. The probability to accept immigrant individuals.

Fitness value	>200	>100	>50	>30	<30
Probability	0.75	0.4	0.1	0.05	0

5.4 CMA-ES with 3 Islands

The 3-island model is equivalent to the pan-population model with the population size of 39 individuals. The strategies used in this variant are the same as for the 6-island model, however, since most of them produces significantly worse results, they are excluded from the plots for the sake of clarity.

5.5 Success Rate Comparison

Figure 3 shows the success rate of examined versions of CMA-ES. Here, the influence of the population size, as well as the island model strategies, are shown. In green, pan-population CMA-ES is shown. There can be seen a strong influence of the population size on the success rate. By optimizing the fitness function in C++, which enabled an increase of the population size up to 78 individuals, the success rate increased from 53% to 87.5%.

Splitting the global population into 6 islands, shown in blue, brings another significant improvement to the success rate. When 3 best individuals are broadcast from each island (6IS3), the success rate reaches 93.3%. To reach a 100% success rate, the acceptance policy has to be turned on. The migration of 3 individuals appears to bring to much genetic material that prevents the CMA-ES to explore local neighborhoods. There are also two cases which are worse than the PP model. In the case of 6ISB, the exchange of genetic material is not sufficient to explore the search space efficiently. This strategy thus reaches about 1% worse success rate. When the acceptance policy is tightened, the success rate of this 6-island CMA-ES drops to 77.4%

Interesting results were achieved with a 3-island version of CMA-ES, shown in yellow. When broadcasting only a single individual and even having a half of the total population, the success rate remains the same as for the 6 islands with

migration of 3 individuals per island. The 3-island model thus seems to be more efficient in terms of computational requirements, which was further confirmed in Fig. 4.

5.6 Number of Evaluations Comparison

Figure 4 presents the difference in the total number of evaluations before the algorithms converge. The simplest 6-island strategies (6IBS and 6IS3) are considerably worse than an equivalent PP version with 78 individuals. Both strategies using the acceptance policy require about the same number of evaluations to converge, however, the superiority of 6IS3AP is emphasized by the best success rate. This version also reaches the most stable performance with only a low variance in the measured number of evaluations.

The PP model with smaller populations needs a significantly lower number of evaluations to converge, but that is redeemed by a poor success rate. The fastest variant uses only 3 islands. Even if the PP model with 40 individuals may sometimes be faster, the 3-island version on average produces the treatment plans in a shorter time, and with much higher quality. The detailed comparison of the best variants can be seen in Fig. 5.

5.7 Acceptance Policy

The acceptance policy defines the probability the immigrants are incorporated into the local population. In this investigation, we calculated the success rate over all islands to see what percentage of the islands could converge to the global optimum and if we can get multiple different optimal solutions. This strategy has shown to be very beneficial for 6-island modes. The highest benefit was

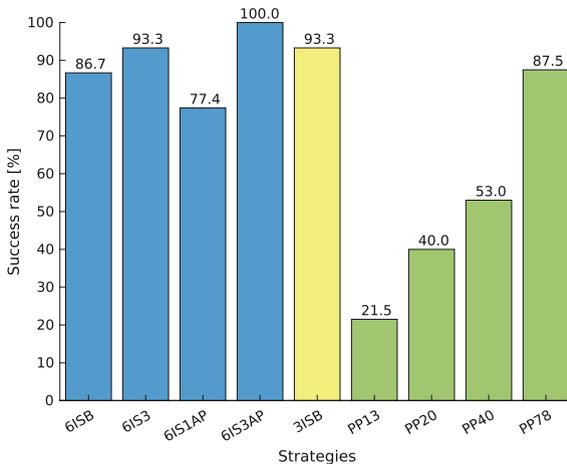


Fig. 3. Success rate of different evolution strategies. (Color figure online)

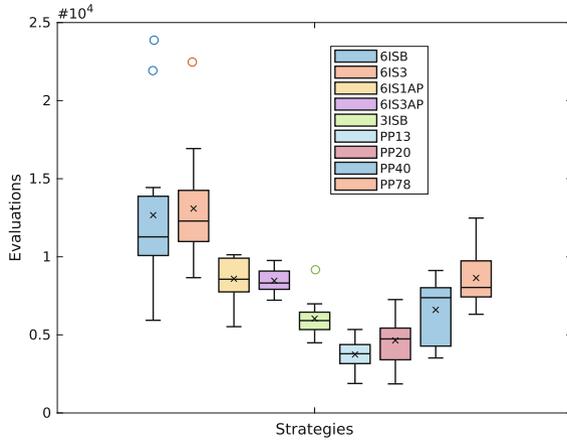


Fig. 4. Number of evaluations before CMA-ES converges.

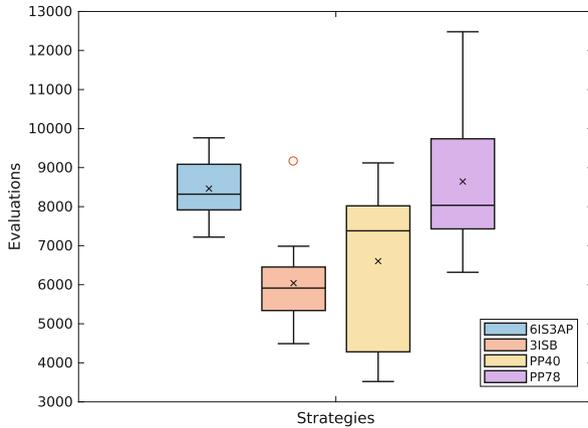


Fig. 5. Comparison of the performance of the best variants of CMA-ES.

brought to S3AP where 2.5 islands out of 6 usually converged to the optimum while the number of evaluations also increased. On the other side, this technique deteriorates the success rate for three-island modes. Most of the trials for 3 islands did not converge to the global optimum, even after a great number of generations, see Table 3.

5.8 Migration Interval

The influence of different migration interval was also tested. The impact of the prolonged migration interval appears to be rather negative. The increase in the migration interval yields an increase in the number of evaluations, and even degradation of the success rate. The modification of migration interval behaves

Table 3. The number of evaluations needed to find one optimal plan and average success rate for different acceptance policies.

Strategy	Evaluations/success rate		
	SB	S1AP	S3AP
6-island	12,900/86.7%	8,430/230%	8,320/252%
3-island	6,220/93.3%	12,500/82.1%	13,600/84.5%

differently for 3-island and 6-island model. The results are shown in Table 4. The reason is much slower convergence which can be clearly seen from Fig. 6. The leftmost plots show a rapid decrease of the best individual fitness values followed by the average and worst individual. Contrary the rightmost plots with prolonged migration interval show much slower convergence.

Table 4. The influence of the migration interval on the average number of evaluations/success rate.

Interval	Number of evaluations/success rate		
	1	2	3
6-island	11,076/86.7%	17,472/73.3%	13,416/86.7%
3-island	5,811/93.3%	8,541/93.3%	7,800/53.3%

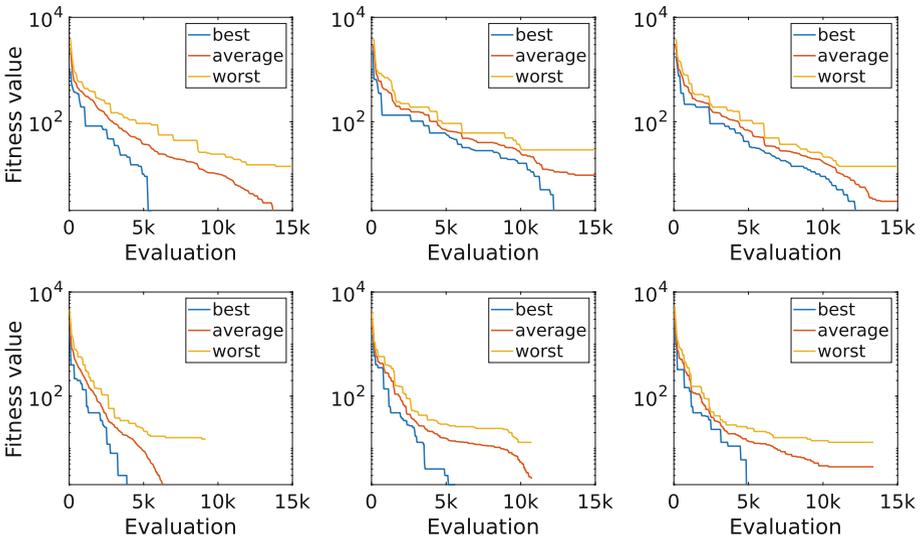


Fig. 6. Convergence of 6ISB on the top and 3ISB strategy on the bottom, for migration interval 1, 2 and 3 from left to right, respectively.

6 Conclusion

This study has presented an acceleration of the CMA-ES algorithm for the design the HIFU treatment plans by the island model. Since the fitness function evaluation took prohibitively long, low-level optimization, parallelization and vectorization in C++ were conducted. This optimization enabled to use twice as big population while fitting into the same maximal time period of 48 h. Next, the effort was put into the parallelization using the island model. Several different parameters on three and six islands were investigated including the number of migrating individuals, the acceptance policy and the length of the migration interval.

The highest success rate was achieved by a 6-island model migrating three best individuals every generation with decreasing acceptance ratio towards the end of the evolution. Comparing the original algorithm with the fastest variant, a 3-island model migrating a single best individual every generation but without the acceptance policy, the evolution was accelerated more than 4 times on the same computer. This allows us to deliver an optimal patient specific treatment plan within 6 h. Practically, the patient can be scanned during the day, and the treatment plan will be computed overnight.

The proposed island model also allows the parallelization of the evolutionary strategy on more interconnected computers, which will be used in the future work to further accelerate the evolution and incorporate a realistic ultrasound model.

Acknowledgment. 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.

References

1. Alba, E., Tomassini, M.: Parallelism and evolutionary algorithms. *IEEE Trans. Evol. Comput.* **6**(5), 443–462 (2002)
2. Cantú-Paz, E.: *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Publishers, Dordrecht (2000)
3. Chen, S., Ma, B., Zhang, K.: On the similarity metric and the distance metric. *Theor. Comput. Sci.* **410**(24–25), 2365–2376 (2009)
4. Cudova, M., Treeby, B.E., Jaros, J.: Design of HIFU treatment plans using an evolutionary strategy. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion on - GECCO 2018*, pp. 1568–1575. ACM Press, New York (2018). <https://doi.org/10.1145/3205651.3208268>
5. Dillon, C., Vyas, U., Payne, A., et al.: An analytical solution for improved HIFU SAR estimation. *Phys. Med. Biol.* **57**(14), 4527–4544 (2012)
6. Frigo, M., Johnson, S.: The Design and Implementation of FFTW3. *Proc. IEEE* **93**(2), 216–231 (2005)

7. Hansen, N.: The CMA evolution strategy: a comparing review. In: *Towards a New Evolutionary Computation*, vol. 192, pp. 75–102. Springer, Heidelberg (2006). https://doi.org/10.1007/3-540-32494-1_4
8. Hansen, N.: *The CMA Evolution Strategy: A Tutorial*. Technical report (apr 2016)
9. Hansen, N., Ostermeier, A.: Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation. In: *Proceedings of IEEE International Conference on Evolutionary Computation*, pp. 312–317. IEEE (1996)
10. Hansen, N., Kern, S.: Evaluating the CMA evolution strategy on multimodal test functions. In: Yao, X., et al. (eds.) *PPSN 2004*. LNCS, vol. 3242, pp. 282–291. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30217-9_29
11. Huijssen, J., Verweij, M.: An iterative method for the computation of nonlinear, wide-angle, pulsed acoustic fields of medical diagnostic transducers. *J. Acoust. Soc. Am.* **127**(1), 33–44 (2010)
12. Ichihara, M., Sasaki, K., Umemura, S., et al.: Blood flow occlusion via ultrasound image-guided high-intensity focused ultrasound and its effect on tissue perfusion. *Ultrasound Med. Biol.* **33**(3), 452–459 (2007)
13. Intel Corporation: *Intel® 64 and IA-32 Architectures Optimization Reference Manual* (2011)
14. Madera, J., Dorronsoro, B.: Estimation of Distribution Algorithms. *Metaheuristic Procedures for Training Neural Networks*, pp. 87–108, December 2006
15. Marquet, F., Perndoiot, M., Aubry, J.F., et al.: Non-invasive transcranial ultrasound therapy based on a 3D CT scan: protocol validation and in vitro results. *Phys. Med. Biol.* **54**(9), 2597–2613 (2009)
16. Massey, J., et al.: AustinMan and AustinWoman: High fidelity, reproducible, and open-source electromagnetic voxel models. In: *34th Meeting of the Bioelectromagnetics Society* (2012)
17. van der Pas, R., Stotzer, E., Terboven, C.: *Using OpenMP-The Next Step: Affinity, Accelerators, Tasking, and SIMD*. MPI Press, Kuala Lumpur (2017)
18. Pennes, H.: Analysis of tissue and arterial blood temperatures in the resting human forearm. *J. Appl. Physiol.* **1**(2), 93–122 (1948)
19. Sapareto, S., Dewey, W.: Thermal dose determination in cancer therapy. *Int. J. Radiat. Oncol. Biol. Phys.* **10**(6), 787–800 (1984)
20. Sudholt, D.: Parallel evolutionary algorithms. In: Kacprzyk, J., Pedrycz, W. (eds.) *Springer Handbook of Computational Intelligence*, pp. 929–959. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-43505-2_46
21. Suomi, V., Jaros, J., Treeby, B., Cleveland, R.: Full modelling of high-intensity focused ultrasound and thermal heating in the kidney using realistic patient models. *IEEE Trans. Biomed. Eng.* **PP**(99), 1–12 (2017)
22. Treeby, B., Jaros, J., Rendell, A., Cox, B.: Modeling nonlinear ultrasound propagation in heterogeneous media with power law absorption using a k-space pseudospectral method. *J. Acoust. Soc. Am.* **131**(6), 4324–36 (2012)
23. Ye, G., Smith, P., Noble, J.: Model-based ultrasound temperature visualization during and following HIFU exposure. *Ultrasound Med. Biol.* **36**(2), 234–249 (2010)
24. Zhou, Y.F., Syed Arbab, A., Xu, R.: High intensity focused ultrasound in clinical tumor ablation. *World J. Clin. Oncol.* **2**(1), 8–27 (2011)