20th International Conference on the Application of Computer
Science and Mathematics in Architecture and Civil Engineering
K. Gürlebeck and T. Lahmer (eds.)
Weimar, Germany, 20-22 July 2015

OPTIMISING ACOUSTIC RESONATORS FOR SONOFUSION EXPERIMENTS WITH EVOLUTIONARY ALGORITHMS

M. J. Stokmaier*, **A. G. Class, T. Schulenberg, and R. T. Lahey Jr.**

**Dynardo GmbH*
Steubenstrasse 25, 99423 Weimar
E-mail: Markus.Stokmaier@dynardo.de

Keywords: evolutionary algorithm, hybrid EA, global optimization, resonator optimization.

Abstract. *The sizing of simple resonators like guitar strings or laser mirrors is directly connected to the wavelength and represents no complex optimisation problem. This is not the case with liquid-filled acoustic resonators of non-trivial geometries, where several masses and stiffnesses of the structure and the fluid have to fit together. This creates a scenario of many competing and interacting resonances varying in relative strength and frequency when design parameters change. Hence, the resonator design involves a parameter-tuning problem with many local optima. As its solution evolutionary algorithms (EA) coupled to a forced-harmonic FE simulation are presented. A new hybrid EA is proposed and compared to two state-of-the-art EAs based on selected test problems. The motivating background is the search for better resonators suitable for sonofusion experiments where extreme states of matter are sought in collapsing cavitation bubbles.*

1 INTRODUCTION

1.1 The optimisation problem of resonator tuning

An important fraction of **parameter tuning problems** in science and engineering is well-behaved in the manner that small variations in the parameter set entail only small changes in the system responses in most parts of the parameter space. The next important question for classifying optimisation (or just function minimisation w.l.o.g.) problems is whether there are local minima apart from the global minimum. Tuning a guitar string is easy because a gradient-following scheme allows to reach the global optimum of matching the frequency target. Minimising the displacement at the middle of the string when it is in forced stationary oscillation by varying tension, cross section, and forcing frequency is also easy as there are many different standing wave patterns with nodes in the centre that can be found by gradient-following. The situation changes, however, when asked to minimise the string displacement at a shifted position, say 2 cm away from the middle. If gaining sufficient problem understanding for an analytical solution is impossible or impractical, then many standing wave patterns have to be tried experimentally or by simulation, i. e. via numerical experiment.

The difficulty level of function minimisation problems depends on (a) the existence of local minima (uni-/multimodal), (b) the number of dimensions, (c) separability (sequential 1D-minimisation along coordinate axes possible?), (d) whether there is structure in the distribution of minima, (e) the existence of cliffs, (f) noise, (g) neutral plateaus, and (h) the relation among attractor sizes (what percentage of the search space funnels into the deepest valley(s)?).

Designing a liquid filled structure which can exhibit a vibration mode allowing the maximisation of the sound pressure amplitude of a standing acoustic wave in the liquid can be easy or hard. In the case of a spherical vessel it is a simple 1D problem of matching the masses and stiffnesses of liquid and solid. But if the geometry becomes more complex, if the number of subsystems in coupled oscillation increases, if the number of masses and stiffnesses multiplies, then the parameter tuning problem becomes harder. Then, simulations for exploring the design space are needed to support the design process, and efficient global optimisation algorithms can be of great help to make this exploration much more goal-oriented.

1.2 Background: the sonofusion experiment

When cavitation bubbles implode, then the bubble's noncondensible gas content is being compressed in the final phase of implosion. At the end, the kinetic energy from the surrounding liquid is transformed into potential energy and concentrated in a tiny amount of gas. If the surface tension manages to conserve the spherical symmetry, then this energy concentration mechanism is capable of raising the energy density by many orders of magnitude above the initial level of the sound wave in the liquid. In fact, the compression-heated gas can turn into light-emitting plasma, an effect which is called *sonoluminescence (SL)*. Where is the limit of this energy concentration mechanism? Could fusion conditions be reached? Could SL become *sonofusion (SF)* [6, 11, 12]? This debated question has too often been asked in the form: "Who will be the first to achieve thermonuclear fusion on the tabletop?". Perhaps it could be more fruitfully put this way: Any SL experiment turns into an SF experiment when deuterium is added to the liquid up front and neutron and tritium detection to the backend. This simply adds one more signal channel to the otherwise sparse data available for characterising SL plasma.

In a cooperation of the Institute of Nuclear and Energy Technologies (IKET) at the Karlsruhe Institute of Technology (KIT) with the Department of Mechanical, Aerospace, and Nuclear Engineering (MANE) at the Rensselaer Polytechnic Institute (RPI) the question has been addressed why the resonators used for the SF experiments of Taleyarkhan et al. [11, 12] are not reproducible with reliable performance and how better resonator designs can be proposed so that the acoustic pressure is maximised and reproducible manufacturing becomes possible. In order to do this, forced-harmonic finite element (FE) analyses of the piezo-driven and liquid-filled acoustic resonators have been subject to parameter tuning with global (GS) and local search (LS) methods like EAs and the downhill-simplex method.

2 EVOLUTIONARY ALGORITHMS (EA) FOR REAL-WORLD PARAMETER TUNING PROBLEMS

When using a blackbox optimiser, a list of tuning parameters $\vec{x} \in \mathbb{R}^n$ (or[1] $\in [0,1]^{n \cdot m}$) is called chromosome or genotype representation of a candidate solution. In real-world applications the phenotype representation of the genotype solution $\vec{x}$ is often a sequence of simulation and postprocessing from which a scalar quality measure $f_{\text{obj}}(\vec{x})$ is distilled. The blackbox optimiser ignores any intermediate phenotype features and merely deals with the response f_{obj}. Deciding to treat a real-world problem with a blackbox optimiser means favouring the advantage of being able to use a generic algorithm of known features over the advantage of being able to leverage detailed phenotype information and human problem understanding during the automated optimisation run. The latter has to be shifted up front (intelligent problem setup and parametrisation for making the task principally (and perhaps efficiently) solvable) which represents an effort, and it may also be used afterwards (learning from solutions found by the optimiser which were undiscovered by previous approaches) which represents another benefit.

2.1 Historic and state-of-the-art EAs

In the 1950s and 60s both things were new and emerging: the knowledge about the substance carrying genetic information in living beings and computer programs made of zeros and ones. It inspired researchers both ways, some wanted to simulate evolution to prove assumptions of how it works in nature, i. e. see how a generation cycle of repetitive random mutation and recombination operations can turn random code of symbols into meaningful code, whereas others wanted to derive efficient optimisation routines for technical real-world problems. Some of these works went far beyond their time in their wide conceptual scope and sophisticated implementations (e. g. [2, 7]). From the 1970s on, two somewhat more narrowly framed EA paradigms, evolution strategies (ES) and genetic algorithms (GA), gained popularity. The good thing about them is that they allow exemplary understanding of several important properties and features of EAs like *gene pool diversity*, *selection pressure*, *exploitation vs. exploration*, or thinking about the geometry of mutation operations, the genotype population, and f_{obj}.

In a classic **genetic algorithm (GA)** the search is coordinate system-dependent because most mutation steps are parallel to coordinate system axes and because component-wise recombinations of two candidate solutions $\vec{x}$ and $\vec{y}$ can only reach the corners of axis-aligned cuboids. GAs stall on diagonal valleys and are often inefficient on nonseparable problems. However, they show the meaning of gene pool diversity and recombination: in many real-world problems not each parameter is heavily coupled with each other parameter; e. g. when subsets of parameters belong to components of a system to be optimised, like the components of a power plant, then it makes sense to try component-wise recombinations of solutions created by re-

[1] in a binary instead of real-coded setup when n parameters are coded with m bits each

combination operators like n-point crossover. The exploration-exploitation conflict appears as the difficulty to choose a good mutation probability: only the mutation operator creates new information, but if its activity is too high, good information is too often destroyed and cannot accumulate in the gene pool.

Classic **evolution strategies (ES)** can teach different lessons: that mutation operators have to be seen as distributions, that the population acts as a low-pass filter scanning for large-scale gradients under small-scale noise, that not all random number distributions retain their shape when the dimension is increased, how self-adaptation can lead to case-sensitive convergence behaviour, and that a local search algorithm can perform a global search when damped self-adaptation leads to the right speed of sinking down from larger to smaller scales in the genotype space. The shortcomings of classic ES are their poor global search power.

Modern EAs are able to lessen or avoid these shortcomings by making creative use of the grown toolbox of mutation and recombination operators and by inventing new ones. In the spirit of [2] they constructively go beyond mere abstractions of nature's EAs. The most performant modern ES, **CMA-ES** [4], adapts its mutation operator in the form of an arbitrarily rotated and stretched multivariate normal distribution (MVN) based on analysing genotype population statistics (instead of mutation & inheritance) and with the result that the population cloud moves and stretches like an amoeba. Another modern EA is **differential evolution (DE)** [10]. It creates a similar tendency of amplifying deviations of the genotype population from the spherical shape by adding vector differences of two population members to a third one. This operation is coordinate system-independent. At the same time, a GA-style component-wise crossover operator acts on each each new chromosome. A third example of the state of the art is **particle swarm optimisation (PSO)** [5] where movements of orbiting particles are modelled through inertia, friction, and attraction. The orbit attractors are stored high-quality solutions once encountered by the swarm (exploitation) while inertia and randomness push exploration. PSO searches more globally when multiple attractors are kept active by not communicating them throughout the full swarm. Similar to the problem of dwindling gene pool diversity in GAs, PSO would have the issue that orbital planes are restricted subspaces of the search space. Therefore, the particle motions are randomised out of the orbital planes.

Hybridising EAs is a way to combine yet more operators and their effects. One goal is to make the algorithm more robust by damping the negative effect if particular operations lead to drag in certain landscapes. Another goal is to increase efficiency, assuming that it can be beneficial if the genetic material produced by one scheme is fed into the next one.

2.2 A simple hybrid EA scheme

The developed in-house hybrid EA [8] can be called tier-based hybrid EA (THEA) as each offspring population is created in five tiers or segments whereby these segments represent ES-, GA-, and DE-like EA paradigms. $\mathcal{P}_g = \{\vec{x}_i\}$ is the evaluated and sorted parent population. The offspring population of same size is initialised and divided into five segments or tiers $\mathcal{P}_{g+1} = \{\mathcal{P}_{elite}, \mathcal{P}_{mutants}, \mathcal{P}_{fp-mutants}, \mathcal{P}_{CO}, \mathcal{P}_{DE}\} = \{\vec{x}'_1, ..., \vec{x}'_{n_1}, \vec{x}'_{n_1+1}, \ldots, \vec{x}'_{n_4}, \vec{x}'_{n_4+1}, ..., \vec{x}'_N\}$, where $n_1 = N_{elite}$, the size of the first tier, $n_2 = N_{elite} + N_{mutants}$ and so on. The offspring generation loop iterates in the following manner: for i in $1, 2, ..., N$:

- **if $\vec{x}'_i$ in 1st tier ("elite"):** Copy $\vec{x}_i$, leave the best chromosome $\vec{x}_1$ untouched, mutate the others weakly (isotropic mutation step distribution, i. e. ES-style).
- **if $\vec{x}'_i$ in 2nd tier ("mutants"):** Copy $\vec{x}_i$ and mutate each vector component with a probability $P_{\text{mut}} < 1$ (GA-style: not all components are mutated).

- **if $\vec{x}'_i$ in 3rd tier ("free parent choice mutants"):** Randomly choose one parent $\vec{x}_k$, copy it, and mutate (GA-style). The random choice is an exponential distribution favouring better parents (GA-style: also the worst can be chosen).

- **if $\vec{x}'_i$ in 4th tier ("CO-bunch"):** Randomly choose (with exponential selection pressure) two numbers $k, l \in [1, N]$, $(k \neq l)$ and form the chromosome $\vec{x}'_i$ by recombining $\vec{x}_k$ with $\vec{x}_l$. This means using a crossover (CO) operator like in genetic algorithms (GA). The two CO operators are uniform CO and line recombination with subsequent isotropic mutation. A GA-style mutation with damped step size follows.

- **if $\vec{x}'_i$ in 5th tier ("DE-bunch"):** Randomly choose three different parents $\vec{x}_k, \vec{x}_l, \vec{x}_o$ to treat them as in DE. $\vec{x}_k$ is chosen with selection pressure, the others without. The difference vector $\vec{x}_o - \vec{x}_l$ is scaled with a random number $\in [0.2, 0.8]$ and added to $\vec{x}_k$. The CO-step is skipped. A GA-style mutation with damped step size follows.

The advantage of the segmented offspring generation is the flexibility to tune the EA blend. One can smoothly fade between extreme cases (e. g. $N_{CO} = N$ is a pure GA) and all kinds of mixtures allowing to either find the most robust blend or to tune to test problems representing a relevant problem class. Since the underlying assumption is that chromosomes generated in one segment are beneficial input for others, only larger population sizes between $40 \leq N \leq 100$ have been used on the resonator tuning task with 10-40 free design parameters. THEA incorporates no self-adaptive features. The only time-varying state variable is a general mutation step size parameter σ which decays exponentially over time. The forced convergence at a fixed rate (as in simulated annealing) is seen as a practical advantage when confronted with a time-consuming simulation to optimise (no risk of premature or delayed convergence).

2.3 Test problems for tuning, benchmarking, and choosing algorithms

Computationally slim test problems enable easy evaluation and benchmarking of blackbox optimisers which is needed during algorithm development and also by prospective applicants for decision making. In order to make a correct choice for the EA tuning of SF resonators, three test problems have been selected as representative and relevant: the Genitor F101 function [13], the CEC-2011 FM-synthesis problem [3], and the charged marble problem [9]. The functions are smooth, multimodal, and they have only weak structures in their distributions of minima. 400 random-initialised optimisation runs were conducted on each problem. Each run included 9600 function calls by the EA and a budget of 400 evaluations for finishing off with local search (downhill-simplex). THEA was compared with CMA-ES and PSO. Table 1 shows the pairwise comparison of the resulting statistics with the Wilcoxon signed ranks test. In agreement with that data, both CMA-ES and THEA but not PSO have been applied to tuning the resonator FE model.

3 OPTIMISED RESONATOR GEOMETRIES

The resonator design used in [11] was rebuilt several times with small variations by the cooperating team at RPI. It is based on manually manufactured glass parts and manual assembly and proved to be very sensitive. Subsequent FEM simulation of this design at IKET and a sensitivity study conducted on the FE model showed that the sensitivity can be explained by the large manufacturing tolerances. The goal of the current project was to propose a resonator redesign with

Table 1: Wilcoxon signed ranks test on EA statistics

The nonparametric Wilcoxon signed ranks test has been applied to statistics of 400 EA+LS runs per algorithm on three selected test problems. The listed p-values indicate the probability of the null hypothesis to be valid, i. e. that the samples come from the same distribution. The comparisons have been declared a tie if $p > 0.1$. The three competing setups of CMA-ES are (A) an own implementation based on Hansen's educational code "barecmaes.py" (`https://www.lri.fr/~hansen/`), (B) his performance code in Python, and (C) the latter with population size-doubling random restarts [1]. The PSO versions are (A) the implementation by Marcel Caraciolo (`https://github.com/marcelcaraciolo/pypso`) with $c_1 = c_2 = 2.05$, and (B) an own implementation with $\alpha = 0.7298$ and $\psi = 2.9922$ and local neighbourhood degree 2. All population sizes were set to 80 except CMA-ES-C starting with 10.

	CMA-ES-A	CMA-ES-B	CMA-ES-C	PSO-A	PSO-B
			charged marbles (8D)		
THEA wins?	win	win	win	win	tie
p-value	7.33e-21	1.66e-14	5.01e-17	1.41e-14	2.39e-01
			FM-synthesis (6D)		
THEA wins?	tie	defeat	win	tie	win
p-value	9.72e-01	5.02e-02	7.80e-05	5.06e-01	4.69e-21
			F101 (10D)		
THEA wins?	win	defeat	win	win	win
p-value	1.80e-05	4.79e-09	1.93e-09	4.70e-33	4.39e-62

much smaller manufacturing tolerances by making precision-machining of the assembly parts possible. The goal is a high-Q system providing a resonance with strong sound pressure amplitude for a given driving voltage while at the same time avoiding elevated sound pressures near fluid-structure interfaces. Only centrally positioned bubble clusters can implode symmetrically. Bubbles in other places destroy the sound field. The function f_{obj} to minimise is thus $-p_{\mathrm{max}}$ on the central axis with a penalty on elevated wall pressure ratios $r_{\mathrm{wp}} = p_{\mathrm{max}}/p_{\mathrm{wall}}$. As f_{obj} is highly multimodal, any new geometry needs to be parametrised and globally optimised in order to be properly evaluated. Figure 1 shows a proposed resonator geometry overcoming many of the shortcomings of the old design. It is displayed in two versions corresponding to two global optimisation results. The first run yielded the design with $p_{\mathrm{max}} = 81$ bar and $r_{\mathrm{wp}} = 0.30$. In order to search for solutions with lower r_{wp} the penalty has been raised for two more optimisation runs. The first was a local search starting from the old solution which couldn't substantially lower r_{wp}. The other run, a new global search, yielded the design depicted in the lower row with $p_{\mathrm{max}} = 32.6$ bar and $r_{\mathrm{wp}} = 0.16$. 32.6 bar Seems quite low in comparison to 81, but it is still 30 % more than the p_{max} of the optimised FE model of the existing resonators (not shown here; driven by just one piezo ring; serving as reference for new designs). There are two limitations: if the p_{max} per driving voltage is too low, then heating due to elevated driving voltages will occur. If r_{wp} is too high, then beyond a certain voltage threshold cavitation on walls will occur while the pressure amplitude in the central region of interest is not yet high enough. Efficient cooling makes it possible to accept lower-Q resonators, drive them with higher voltage, and benefit from a low-r_{wp} layout. A renewed dedicated campaign of experimental resonator trials is necessary to explore the tradeoffs between p_{max}, r_{wp}, and heating. Simulation-based optimisation techniques help understand the tradeoffs and guide experimental work.

4 CONCLUSION

Designing new resonators for SF experiments involves a hard parameter tuning problem. Comparing design points without pre-optimisation makes little sense. With local search techniques performant resonator setups can only be found with a lucky starting point. Yet, with

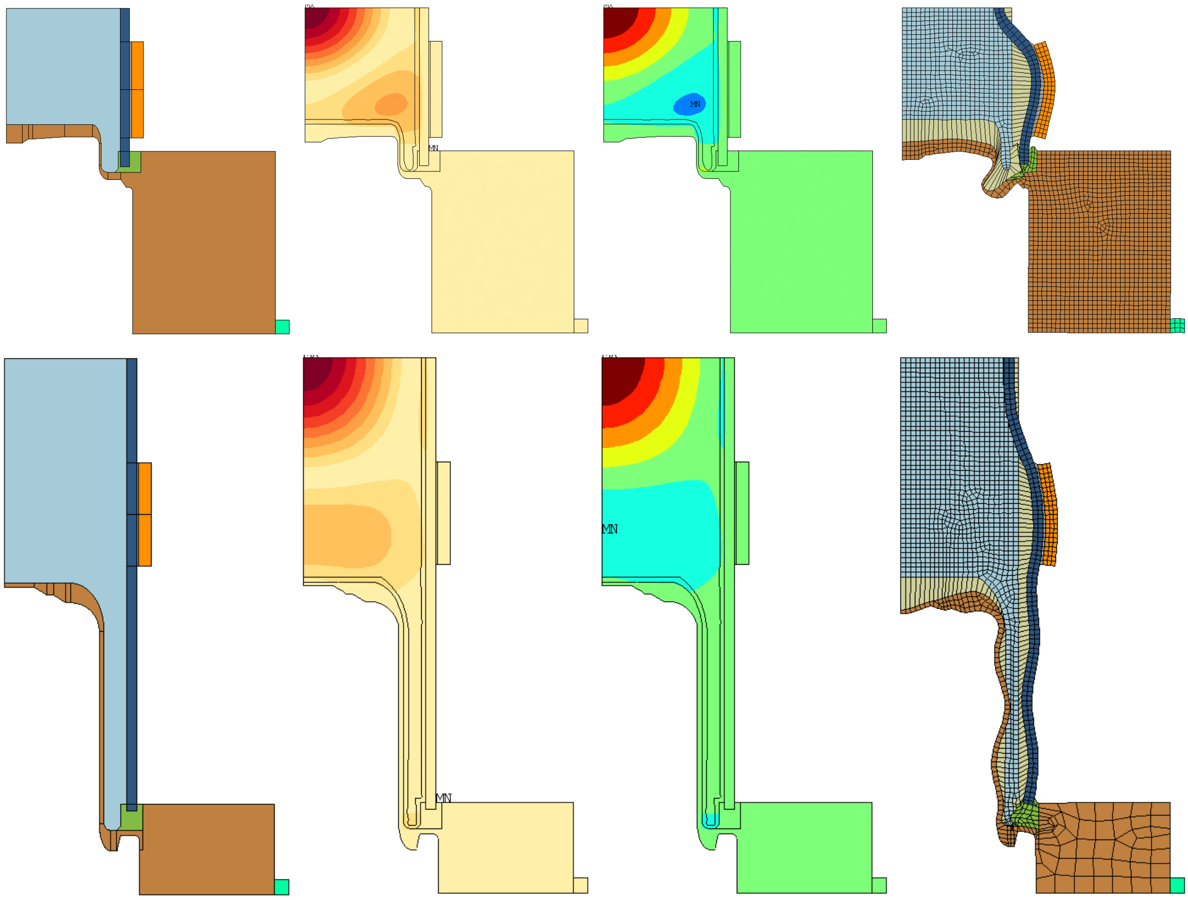

Figure 1: Optimised resonator geometry
Two optimised resonator designs are shown. As the FE model (Ansys® Mechanical APDL) is 2D-axisymmetric and has a horizontal symmetry plane, only a quarter of the cross section is shown. The materials are glass (blue), aluminium (brown), acetone (light blue), piezo driver (orange), silicone (green), soft fixation (cyan). The four picture types from left to right are: geometry, pressure amplitude, pressure snapshot (red – pressure, blue – tension), and deformed mesh (displacement scaling = 800). Both designs are the results of random-initialised global searches with THEA and subsequent downhill-simplex local search on the same base geometry and parametrisation. The only difference is the threshold setting for where the smooth wall pressure ratio-based penalty function sets in. The EA run leading to the upper design was with $N = 80$ over $G = 30$ generations followed by 960 iterations of the downhill-simplex (DS) algorithm. The second EA run was over $G = 39$ generations and 600 DS iterations. The scanned frequency intervals were 16-24 kHz in the first and 18-28 kHz in the second case for the EA search; LS was allowed to shift the range. The found working points are at $f = 25.9$ kHz, $p_{\max} = 81$ bar, $r_{\text{wp}} = 0.30$ and $f = 20.9$ kHz, $p_{\max} = 32.6$ bar, $r_{\text{wp}} = 0.16$, respectively. The higher wall pressure ratio of the upper design is mirrored by the existence of the elliptical dark blue region in the pressure snapshot. It has the same close distance from the surfaces of piston and glass wall. In the course of the local search the maximum pressure on the two surfaces becomes equal. Since LS starting on the upper design with the sharpened penalty could not reduce r_{wp} and since the resonator had to become longer to host more antinodes along the central axis, it can be inferred that it was necessary to give the sound field more room to decay towards the pistons and that the change in the mode shape ivolves surmounting a hill in f_{obj}.

an efficient global optimiser at hand the iterative development cycle towards deeper problem understanding and better resonator designs becomes quicker and much more systematic. EA hybridisation is a way to incorporate and combine lessons learnt from historic paradigm-driven EA approaches. Segmentation of the offspring population is a simple hybridisation approach offering the advantage of flexible EA blending ratios.

REFERENCES

[1] A. Auger, N. Hansen: A restart CMA evolution strategy with increasing population size: Proc. IEEE Congress on Evolutionary Computation (2005), pp. 1769–1776.

[2] H. J. Bremermann: Numerical optimization procedures derived from biological evolution processes. Cybernetic Problems in Bionics (1968), pp. 597–615.

[3] S. Das, P. N. Suganthan: Problem Definitions and Evaluation Criteria for the CEC 2011 Competition on Testing Evolutionary Algorithms on Real-World Optimization Problems. Tech. Rep., Jadavpur University (2010).

[4] N. Hansen, A. Ostermeier: Completely Derandomized Self-Adaptation in Evolution Strategies. Evolutionary Computation **9** (2001), pp. 159–195.

[5] J. Kennedy, R. Eberhart: Particle swarm optimization. Proc. IEEE International Conference on Neural Networks (1995), pp. 1942–1948.

[6] R. I. Nigmatulin, I. Sh. Akhatov, A. S. Topolnikov, R. Kh. Bolotnova, N. K. Vakhitova, R. T. Lahey Jr., R. P. Taleyarkhan: Theory of supercompression of vapor bubbles and nanoscale thermonuclear fusion. Physics of Fluids **17** (2005), p. 107106.

[7] J. Reed, R. Toombs, N. A. Baricelli: Simulation of biological evolution and machine learning: I. Selection of self-reproducing numeric patterns by data processing machines, effects of hereditary control, mutation type and crossing. Journal of Theoretical Biology **17** (1967), pp. 319–342.

[8] M. J. Stokmaier, A. G. Class, T. Schulenberg, R. T. Lahey Jr.: Sonofusion: EA optimisation of acoustic resonator. PAMM **12-1** (2012), pp. 623-624.

[9] M. J. Stokmaier, A. G. Class, T. Schulenberg: A hard optimisation test function with symbolic solution visualisation for fast interpretation by the human eye. Proc. IEEE Congress on Evolutionary Computation (2013), pp. 2251–2258.

[10] R. Storn, K. Price: Differential Evolution A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. Journal of Global Optimization **11-4** (1997), pp. 341–359.

[11] R. P. Taleyarkhan, C. D. West, J. S. Cho, R. T. Lahey Jr., R. I. Nigmatulin, R. C. Block: Evidence for Nuclear Emissions During Acoustic Cavitation. Science **295** (2002), p. 1868.

[12] R. P. Taleyarkhan, J. S. Cho, C. D. West, R. T. Lahey Jr., R. I. Nigmatulin, R. C. Block: Additional Evidence for Nuclear Emissions During Acoustic Cavitation. Phys. Rev. E **69** (2004), p. 036109.

[13] D. Whitley, S. Rana, J. Dzubera, K. E. Mathias: Evaluating evolutionary algorithms. Artificial Intelligence **85** (1996), pp. 245–276.