A TUNABLE REAL-WORLD MULTI-FUNNEL BENCHMARK PROBLEM FOR EVOLUTIONARY OPTIMIZATION

And Why Parallel Island Models Might Remedy the Failure of CMA-ES on It

Christian L. Müller, Ivo F. Sbalzarini
Institute of Theoretical Computer Science and Swiss Institute of Bioinformatics,
ETH Zurich, Universitätstrasse 6, 8092 Zurich, Switzerland
christian.mueller@inf.ethz.ch, ivos@ethz.ch

Keywords: Benchmark, Multi-funnel landscape, Lennard-Jones cluster, CMA-ES, Parallel island model.

Abstract: A common shortcoming in the Evolutionary Computation (EC) community is that the publication of many search heuristics is not accompanied by rigorous benchmarks on a balanced set of test problems. A welcome effort to promote such test suites are the IEEE CEC competitions on real-valued black-box optimization. These competitions prescribe carefully designed synthetic test functions and benchmarking protocols. They do, however, not contain tunable real-world examples of the important class of multi-funnel functions. We argue that finding minimum-energy configurations of 38-atom Lennard-Jones (LJ_{38}) clusters could serve as such a benchmark for real-valued, single-objective evolutionary optimization. We thus suggest that this problem be included in EC studies whenever general-purpose optimizers are proposed. The problem is tunable from a single-funnel to a double-funnel topology. We show that the winner of the CEC 2005 competition, the Evolution Strategy with Covariance Matrix Adaptation (CMA-ES), works on the single-funnel version of this test case, but fails on the double-funnel version. We further argue that this performance loss of CMA-ES can be relaxed by using parallel island models. We support this hypothesis by simulation results of a parallel island CMA-ES, the Particle Swarm CMA-ES, on a subset of the multi-funnel functions in the CEC 2005 benchmark.

1 INTRODUCTION

High-dimensional, non-convex, and noisy optimization problems are commonplace in physics, engineering, and biology. The prominent example of protein structure prediction can, e.g., be considered a minimization problem over an empirical energy landscape. In many cases, such problems can only be tackled efficiently by evolutionary optimization methods such as the Evolution Strategy with Covariance Matrix Adaptation (CMA-ES) (Hansen and Ostermeier, 2001), a particularly successful gradient-free heuristic for real-valued black-box optimization. During exploration of the search space, CMA-ES samples a population of candidate solutions from a multivariate Gaussian distribution. The mean and the covariance matrix of the sampling distribution are continuously adapted in order to guide the search toward promising regions of the space.

Assessing the efficacy and efficiency of evolutionary algorithms is empirical by nature since theoretical convergence analyses are missing for most problems of practical importance. Unfortunately, it is a common shortcoming in the Evolutionary Computation (EC) community that many publications on new search heuristics are not accompanied by rigorous benchmarks on a diverse and balanced set of test problems. It is still commonplace that authors choose an arbitrary set of popular benchmark functions and compare the performance of their algorithm to that of an arbitrary set of existing heuristics. A welcome exception is the series of IEEE CEC competitions on real-valued black-box optimization. Starting in 2005 with a contest on single-objective black-box optimization (Suganthan et al., 2005), these competitions prescribe carefully designed *synthetic* function sets and benchmarking protocols. A generally accepted standard set of *real-world* benchmark problems, however, is still not available.

While studies on synthetic test suites allow fair and re-

producible comparison between algorithms, they only provide limited insight into *why* a particular heuristic works on certain functions, but fails on others. In recent years, it has been suggested that the global topology of a problem may have a strong influence on the performance of search heuristics. (Hansen and Kern, 2004) and (Sutton et al., 2006), for example, suggested that the performance of CMA-ES and Particle Swarm Optimization (PSO) can strongly decrease on test functions with a *multi-funnel* topology, where local optima cannot be interpreted as perturbations to an underlying convex (unimodal) topology. (Lunacek et al., 2008) further examined this issue by considering the search performance of CMA-ES on synthetic, tunable double-funnel landscapes. The global structure of these landscapes consists of a composition of two sphere functions whose sizes and depths are adjustable. They report two remarkable observations. First, the probability of finding the minimum with CMA-ES *without restarts* decreases with increasing population size when considering a double-sphere landscape where the global minimum lies in the smaller funnel. Second, introduction of local multi-modality in the double-sphere function through a Rastrigin function triggers an irresolvable trade-off for the optimal population size of CMA-ES: Local multi-modality demands larger populations while the double-funnel favors the opposite. Therefore, the probability of finding the global minimum with CMA-ES on the designed double-Rastrigin function is low for any population size. Unfortunately, problems with a multi-funnel structure are not of mere academic interest. Prominent real-world examples include trajectory planning in space missions, protein aggregation and misfolding (Clark, 2004), and the potential-energy surfaces (PES) of atomic clusters (Wales, 2004).

The purpose of the present paper is twofold. First, we propose a tunable real-world benchmark problem for real-valued, single-objective, gradient-free evolutionary optimization: the prediction of minimum-energy configurations of the 38-atom Lennard-Jones (LJ_{38}) cluster. This problem has been intensively studied in physical chemistry, and we expect this to be valuable to the EC community. The LJ_{38} problem exhibits a double-funnel PES topology (Doye et al., 1999), which can be transformed to a single-funnel topology by adding a compression penalty. This transformation is tunable by a single, scalar parameter. We show that the LJ_{38} cluster problem cannot be solved by the winner of IEEE CEC 2005 contest, a restart variant of CMA-ES with iteratively increasing population size (IPOP-CMA-ES) (Auger and Hansen, 2005), unless the compression penalty is included. We therefore propose that the LJ_{38} problem should be included in future EC studies whenever a novel general-purpose optimizer is introduced.
Second, we propose that the use of parallel island models can remedy the decreased performance of IPOP-CMA-ES on multi-funnel functions. We believe that it is beneficial to evolve several communicating CMA-ES instances on the problem in parallel, rather than increasing the population size of a single CMA-ES instance. We support this hypothesis by simulation results of a parallel island CMA-ES, the Particle Swarm CMA-ES (PS-CMA-ES), on a subset of multi-funnel functions from the IEEE CEC 2005 test suite.
This paper is organized as follows: In Section 2 we briefly review the IEEE CEC 2005 benchmark test suite and introduce the real-world multi-funnel energy landscape of the 38-atom Lennard-Jones cluster. Section 3 reviews (IPOP-)CMA-ES, the concept of parallel island models, and PS-CMA-ES. In Section 4 we present the results of numerical experiments with IPOP-CMA-ES and PS-CMA-ES. Section 5 discusses these results and concludes this work.

2 MULTI-FUNNEL OPTIMIZATION PROBLEMS

Problems with a multi-funnel topology arise in diverse scientific areas such as, e.g., trajectory planning in space missions or protein aggregation and misfolding in biology. Although there is no clear mathematical definition of the term *multi-funnel*, we consider a problem multi-funnel if it has no underlying convex topology on a global length scale. Probably the best-known example for a multi-funnel function is Schwefel's problem. The Rastrigin function, on the other hand, is highly multi-modal, but not multi-funnel due to its globally spherical topology. In this section, we first revisit the synthetic multi-funnel problems of the IEEE CEC 2005 test suite and then propose the LJ_{38} problem as a tunable real-world test case for gradient-free optimization.

2.1 The IEEE CEC 2005 Test Suite

One key attempt to provide a standard for performance evaluation and analysis of real-valued search heuristics in the EC community is the IEEE CEC 2005 test suite (Suganthan et al., 2005). The benchmark suite includes a set of 25 test functions from a variety of classes such as uni-/multimodal, (non-) separable, noisy, (a-)symmetric, and scalable. Functions f11 to f13 and f15 to f25 are multi-funnel. Along with

the test functions, the suite also specifies a detailed protocol how to evaluate a search heuristic: the problem dimension ranges from $n = 10 \ldots 50$, the number of allowed function evaluations is restricted to $n \cdot 10^4$, measures for success performance are defined, and examples of how to present the results in tables and figures are provided. We refer to the original publication for the full description of the test suite (Suganthan et al., 2005). Here, we use functions f11 to f13, f15, and f16 to evaluate the performance of IPOP- and PS-CMA-ES on multi-funnel problems.

2.2 Lennard-Jones Clusters

The PES of Lennard-Jones (LJ) clusters is one of the best studied models both in physical chemistry and, to some extent, in the field of global optimization. The objective of finding the minimum-energy configuration of this model has fascinated researchers from both communities for over three decades. Each pair of atoms in a LJ cluster interacts through the LJ pair potential

$$u(r_{ij}) = 4\varepsilon \left(\left(\frac{\sigma_{\mathrm{LJ}}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{\mathrm{LJ}}}{r_{ij}} \right)^{6} \right), \qquad (1)$$

where r_{ij} is the distance between atoms X_i and X_j, ε the potential-well depth (in units of energy) and $2^{\frac{1}{6}}\sigma_{\mathrm{LJ}}$ the equilibrium inter-atom distance (in units of length) at zero temperature. The potential energy E_{LJ} of a cluster of N LJ atoms is given by:

$$E_{\mathrm{LJ}} = 4\varepsilon \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \left(\left(\frac{\sigma_{\mathrm{LJ}}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{\mathrm{LJ}}}{r_{ij}} \right)^{6} \right). \qquad (2)$$

In chemistry, the LJ potential is widely used to model hydrophobic interactions in polymers or the behavior of (noble) gases such as Argon. Knowledge about minimum-energy (or ground-state) configurations of LJ clusters enables, e.g., predicting properties of crystallization or solid–liquid transitions at low temperatures. At the same time, LJ clusters are a model problem for general-purpose non-convex optimization. Minimizing the potential energy of a cluster of N atoms in 3D space defines a continuous optimization problem in $3N-6$ dimensions since 3 translational and 3 rotational degrees of freedom can be removed from the system. This is achieved by placing the first atom at the origin of the Cartesian coordinate system, the second along the x-axis and the third in the $x--y$ plane. Several general-purpose optimizers have been tested on LJ clusters of various sizes, most notably the Basin Hopping algorithm (Wales and Doye, 1997), which makes extensive use of gradient information. To date, finding the minimum-energy configuration of LJ clusters is only feasible for global search heuristics that are allowed to exploit exact, analytical gradients. All known putative minimum-energy configurations are listed in the Cambridge Cluster Database, providing a valuable resource for optimization benchmarks. A particularly well-understood configuration is the 38-atom cluster, forming the basis for the new benchmark problem proposed here.

2.2.1 The Energy Landscape Of The 38-atom Lennard-Jones Cluster

The topology of the LJ_{38} PES has been extensively studied in the literature (Doye et al., 1999; Doye, 2000; Wales, 2004). It exhibits a double-funnel structure where the global minimum-energy configuration with face-centered cubic octahedral (fcc) symmetry lies in a narrow funnel and the majority of local minima, all with icosahedral (ico) symmetry, populate the wider funnel. Figure 1 shows a sketch of the landscape.

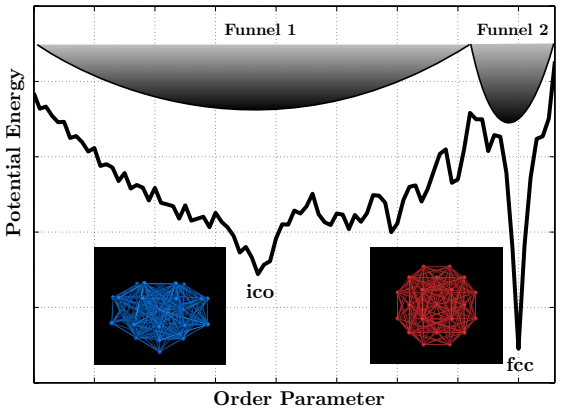

Figure 1: Sketch of the PES of the 38-atom LJ cluster. The x-axis represents a suitable order parameter that can discriminate between different cluster topologies, the y-axis represents the potential energy of the configurations. The wide funnel on the left is populated by clusters with icosahedral (ico) symmetry. The lowest-energy configuration with ico symmetry is depicted in blue. The narrow funnel on the right contains the global minimum with face-centered cubic octahedral (fcc) symmetry (depicted in red).

(Doye et al., 1999) characterized this landscape in great detail in terms of the number and locations of minima, structural diversity of the minima, and the energy barrier between the two funnels. In a subsequent study, (Doye, 2000) also studied the effect of compression on the PES of the 38-atom LJ cluster. Compression can be seen as a transformation of the PES that favors more compact structures. Hence, it broadens the funnel of the more spherical fcc structures and narrows the ico funnel. Doye added the fol-

lowing compression penalty to the energy function:

$$E_{\mathrm{CLJ}} = E_{\mathrm{LJ}} + \mu_{\mathrm{comp}} \sum_{i=1}^{N} \frac{\|X_i - X_{\mathrm{cm}}\|^2}{\sigma_{\mathrm{LJ}}^2}, \qquad (3)$$

where X_i is the position of atom i and X_{cm} is the center of mass of the cluster. The parameter μ_{comp} controls the magnitude of the compression acting on the configuration. When $\mu_{\mathrm{comp}} \geq 5\varepsilon$, the PES of the cluster exhibits a clear single-funnel topology. The effect of the μ_{comp}-dependent compression on the PES is beautifully visualized in (Wales, 2004), pp. 338–339, and (Doye, 2000). Note that even with the compressed potential E_{CLJ} the system contains a staggering number of local minima, but the global topology can be changed from double-funnel to single-funnel. This renders the 38-atom LJ cluster with tunable compression an ideal real-world test case to study the performance of evolutionary algorithms and their sensitivity to the underlying topology.

2.2.2 Bond-order Parameters For Cluster Characterization

Bond-order parameters are indispensable for the analysis of cluster configurations (Steinhardt et al., 1983). These parameters, termed Q_4, Q_6, W_4, and W_6 are able to discriminate symmetry groups in crystals and liquids. The parameters Q_4 and Q_6 have been identified as the most informative ones for distinguishing fcc from ico configurations. Q_l is defined as:

$$Q_l = \left(\frac{4\pi}{2l+1} \sum_{m=-l}^{l} \|\bar{Q}_{lm}\|^2 \right)^{\frac{1}{2}}, \qquad (4)$$

where

$$\bar{Q}_{lm} = \frac{1}{N_b} \sum_{r_{ij} < r_0} Y_{lm}(\Theta_{ij}, \Phi_{ij}). \qquad (5)$$

Here, N_b denotes the number of pseudo bonds that have a length smaller than a cutoff distance r_0. $Y_{lm}(\Theta_{ij}, \Phi_{ij})$ are spherical harmonics with Θ_{ij} being the polar and Φ_{ij} the azimuthal angle of the interatomic vector $\mathbf{r}_{ij}$ with respect to an arbitrary coordinate frame. We propose to use the order parameter Q_4 with a cutoff of $r_0 = 1.391\sigma_{\mathrm{LJ}}$ as suggested for LJ cluster analysis (Doye et al., 1999). A Q_4 value of 0 corresponds to a perfect icosahedron, a value of 0.1909 to a perfect fcc configuration.

3 STANDARD CMA-ES AND PARALLEL-ISLAND CMA-ES

This section reviews (IPOP-)CMA-ES and parallel island models for CMA-ES.

3.1 The CMA Evolution Strategy

The standard CMA-ES uses weighted intermediate recombination, step size adaptation, and a combination of rank-μ update and rank-one update (Hansen and Ostermeier, 2001; Hansen, 2007). At each iteration of the algorithm, the members of the new population are sampled from a multivariate normal distribution $\mathcal{N}$ with mean $\mathbf{m} \in \mathbb{R}^n$ and covariance $\mathbf{C} \in \mathbb{R}^{n \times n}$. The sampling radius is controlled by the overall standard deviation (step size) σ. Let $\mathbf{x}_k^{(g)}$ the k^{th} individual at generation g. The new individuals at generation $g+1$ are sampled as:

$$\mathbf{x}_k^{(g+1)} \sim \mathbf{m}^{(g)} + \sigma^{(g)} \mathcal{N}\left(\mathbf{0}, \mathbf{C}^{(g)}\right) \quad k = 1, \ldots, \lambda. \qquad (6)$$

The λ sampled points are then ranked in order of ascending fitness, and the μ best are selected. The mean of the sampling distribution is updated using weighted intermediate recombination (Hansen, 2007). The covariance matrix is adapted as described in various publications and reports (Hansen and Ostermeier, 2001; Hansen, 2007). The behavior of CMA-ES is mainly controlled by two parameters: the initial step size σ and the population size λ. In standard CMA-ES, the population size is chosen as $\lambda = 4 + \lfloor 3 \ln n \rfloor$ (Hansen and Ostermeier, 2001). The most successful variant of CMA-ES, IPOP-CMA-ES, employs a restart mechanism with iterative doubling of the population size until a stopping criterion is met (Auger and Hansen, 2005). It uses an initial step size of $\sigma = (B - A)/2$, where $[A, B]^n$ is the bounded search space. These standard choices render IPOP-CMA-ES quasi parameter-free.

3.2 Parallel Island Models for CMA-ES

The concept of parallel island models for evolutionary algorithms is to use a set of independent (possibly heterogeneous) populations (islands) that concurrently explore the problem space. During search, the populations are allowed to migrate (communicate information) in an (a-)synchronous way. Many different migration patterns and modes of information sharing have been examined in recent years. Two advantages of parallel island models are that (a) different distant parts of the search space can be explored simultaneously and integrated into individual populations, and (b) such algorithms can easily be parallelized in a distributed computing environment. For detailed information on the topic, we refer to (Alba, 2005).

A specific instance of a parallel island model for CMA-ES, the Particle Swarm CMA-ES (Müller et al., 2009), has been introduced in order to improve the performance of CMA-ES on multi-funnel landscapes.

Inspired by ideas from Particle Swarm Optimization (PSO), the PS-CMA-ES algorithm evolves a swarm of S CMA-ES instances *in parallel*. The individual instances exchange promising solutions every I_c generations. Both the swarm size S (the total number of concurrent CMA-ES instances) and the communication interval I_c are strategy parameters. The influence of global swarm information on the covariance matrix adaptation of each CMA-ES instance is controlled by a weight parameter $c_p \in [0,1]$. A detailed description of the algorithm can be found in (Müller et al., 2009).

4 COMPUTATIONAL EXPERIMENTS

We first show the performance of IPOP-CMA-ES on LJ_{38} and then assess IPOP- and PS-CMA-ES on a multi-funnel subset of the CEC 2005 test suite.

4.1 The 38-Atom LJ Cluster

We search for the minimum-energy configurations of LJ_{38} in a 3D $[-4,4]$ box. The search space thus is $[-4,4]^n$ with dimension $n = 3 \cdot 38 - 6 = 108$. The LJ parameters σ_{LJ} and ε are both set to 1. In order to assess the effectiveness of IPOP-CMA-ES on LJ_{38}, we follow the parameterization and boundary handling of (Auger and Hansen, 2005) with two important exceptions. Instead of doubling the population size after each restart, we increase λ by a factor of 1.25. We also refine the initialization procedure in order to assess the effect of initial LJ-particle positions on the quality of the solution: We place the initial population mean uniformly random in the box $[-0.5,0.5]^3$, $[-1.5,1.5]^3$, or $[-3,3]^3$ and sample the initial population with a σ of 20% of the box edge length. We repeat the experiment 25 times for each box size with no bounds on the number of allowed function evaluations. We stop an optimization run whenever the population size exceeds the initial λ by a factor of 100, i.e., after 21 restarts. We consider LJ_{38} clusters without compression and with $\mu_{comp} = 5\varepsilon$.

Figure 2 summarizes the results. The experiments reveal that on the LJ_{38} problem without compression none of the IPOP-CMA-ES runs reach the lowest-energy fcc configuration, nor the lowest-energy ico configuration (blue □ in Figure 2), despite the fact that each run used, on average, 10^8 function evaluations. With compression, however, *all* IPOP-CMA-ES runs find the globally optimal fcc configuration, independent of initialization. Even more strikingly, all local minima found on the problem without compression have icosahedral-like configurations

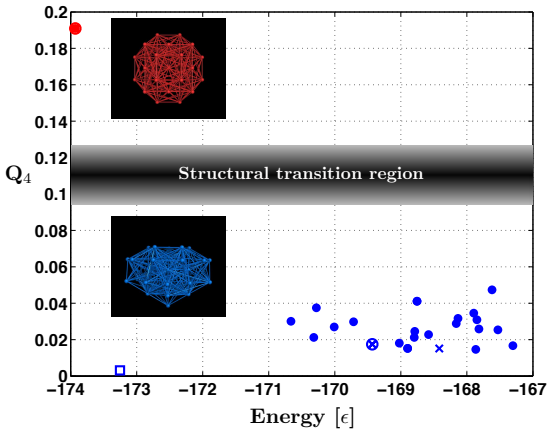

Figure 2: Bond-order parameter Q_4 vs. potential energy (in units of ε) for all minima found by IPOP-CMA-ES. The blue data points show the $3 \cdot 25$ local minima found on the LJ_{38} problem without compression with the initial $[-0.5,0.5]^3$ box ($\times$), the $[-1.5,1.5]^3$ box ($\circ$), and the $[-3,3]^3$ box ($\bullet$). The single blue □ in the bottom-left corner marks the lowest-energy icosahedral configuration. The shaded gray area is the structural transition region from ico to fcc symmetry (Doye et al., 1999). The red data point in the upper-left corner corresponds to the global minimum ($E_{LJ} = -173.9284\varepsilon$), which was found by IPOP-CMA-ES in all $3 \cdot 25$ runs when using the compression penalty.

with Q_4 values between 0.01 and 0.045. The runs that started in the largest box (blue $\bullet$'s in Fig. 2) show the largest structural diversity of the minima ($Q_4 = 0.01 \ldots 0.045$), including the lowest local minimum found at $E_{LJ} \approx -170.66\varepsilon$. All runs that started in the $[-1.5,1.5]^3$ box converged to the same local minimum, while the runs in the smallest box found an additional local minimum.

4.2 The IEEE CEC 2005 Test Suite

We provide evidence that the failure of IPOP-CMA-ES on multi-funnel functions can be remedied by using a parallel island model instead of a single population of increasing size. Therefore, we compare the performance of the parallel island PS-CMA-ES to that of IPOP-CMA-ES on the multi-funnel functions f11 to f13, f15, and f16 of the CEC 2005 benchmark suite. The results are summarized in Table 1. All experiments are conducted according to the CEC 2005 protocol. While IPOP-CMA-ES is still quite efficient on the modified Weierstrass function (f11) and Schwefel's problem (f12), it is not able to solve any of the other multi-funnel functions in the set. PS-CMA-ES can solve f11 to f13 with strategy parameters $c_p = 0.7$, $I_c = 200$, and $S = 15$, and f15 and f16 with strategy parameters $c_p = 0.3$, $I_c = 150$, and $S = 10$. To our knowledge, no other algorithm has so far been re-

ported to solve the pair f15/f16.

Table 1: Summary statistics of IPOP-CMA-ES (first block, taken from (Auger and Hansen, 2005)) and PS-CMA-ES (second block) on the multi-funnel functions f11 to f13, f15, and f16 of the CEC 2005 test suite in $n = 10$ dimensions. The minimum, median, maximum, and mean number of function evaluations needed to solve the problems are reported over 25 repetitions of each test. The last column reports the success rates. See main text for strategy parameters used.

	min	med	max	mean	Suc. Rate
f11	3.05e+04	-	-	6.31e+4	2.40e-01
f12	2.37e+03	3.10e+04	-	2.88e+04	8.80e-01
f13	-	-	-	-	0
f15	-	-	-	-	0
f16	-	-	-	-	0
f11	3.56e+04	-	-	5.02e+04	4.00e-01
f12	1.66e+04	1.89e+04	2.13e+04	1.92e+04	1.00e+00
f13	8.34e+04	-	-	8.34e+04	4.00e-02
f15	3.30e+04	-	-	3.65e+04	1.20e-01
f16	5.09e+04	-	-	5.09e+04	4.00e-02

5 CONCLUSIONS

In summary, we have proposed: (a) a tunable real-world multi-funnel test case for gradient-free evolutionary optimizers and (b) that parallel island models can relax the performance decrease of CMA-ES on multi-funnel functions where the global optimum is not located in the broadest funnel. The latter can be explained by the irresolvable trade-off for the optimal population size of CMA-ES on multi-funnel functions (Lunacek et al., 2008). The larger the population size, the more likely CMA-ES converges in the broadest, sub-optimal funnel. This is supported by the fact that IPOP-CMA-ES performs remarkably well when the PES of LJ_{38} is compressed to a single-funnel structure and that the parallel island PS-CMA-ES outperforms IPOP-CMA-ES on the considered sub-set of multi-funnel functions from the CEC 2005 test suite. To date, no population-based, gradient-free evolutionary algorithm can solve the presented LJ_{38} test case without compression. Since many real-world applications entail multi-funnel landscapes, we argue that the EC community should focus on algorithms that can solve such problems. As a benchmark, we suggest the LJ_{38} problem, which should be considered in all future EC studies that introduce a novel general-purpose optimizer.

ACKNOWLEDGMENTS

We thank Georg Ofenbeck for setting up the LJ cluster simulations.

REFERENCES

Alba, E. (2005). *Parallel Metaheuristics: A New Class of Algorithms*. Wiley-Interscience.

Auger, A. and Hansen, N. (2005). A restart CMA evolution strategy with increasing population size. In *Proc. IEEE Congress on Evolutionary Computation (CEC 2005)*, volume 2, pages 1769–1776.

Clark, P. L. (2004). Protein folding in the cell: reshaping the folding funnel. *Trends Biochem. Sci.*, 29(10):527–534.

Doye, J. (2000). Effect of compression on the global optimization of atomic clusters. *Phys. Rev. E*, 62(6, Part B):8753–8761.

Doye, J. P. K., Miller, M. A., and Wales, D. J. (1999). The double-funnel energy landscape of the 38-atom Lennard-Jones cluster. *J. Chem. Phys.*, 110(14):6896–6906.

Hansen, N. (2007). The CMA Evolution Strategy: A Tutorial. http://www.lri.fr/ hansen/cmatutorial.pdf.

Hansen, N. and Kern, S. (2004). Evaluating the CMA Evolution Strategy on Multimodal Test Functions. In *Lect. Notes Comput. Sc.*, volume 3242 of *Parallel Problem Solving from Nature - PPSN VIII*, pages 282–291.

Hansen, N. and Ostermeier, A. (2001). Completely Derandomized Self-Adaption in Evolution Strategies. *Evolutionary Computation*, 9(2):159–195.

Lunacek, M., Whitley, D., and Sutton, A. (2008). The Impact of Global Structure on Search. In *Lect. Notes Comput. Sc.*, volume 5199 of *Parallel Problem Solving from Nature - PPSN X*, pages 498–507.

Müller, C. L., Baumgartner, B., and Sbalzarini, I. F. (2009). Particle swarm CMA evolution strategy for the optimization of multi-funnel landscapes. In *Proc. IEEE Congress on Evolutionary Computation (CEC 2009)*, pages 2685–2692.

Steinhardt, P. J., Nelson, D. R., and Ronchetti, M. (1983). Bond-orientational order in liquids and glasses. *Phys. Rev. B*, 28(2):784–805.

Suganthan, P. N., Hansen, N., Liang, J. J., Deb, K., Chen, Y.-P., Auger, A., and Tiwari, S. (2005). Problem Definitions and Evaluation Criteria for the CEC 2005 Special Session on Real-Parameter Optimization. Technical report, Nanyang Technological University, Singapore.

Sutton, A. M., Whitley, D., Lunacek, M., and Howe, A. (2006). PSO and multi-funnel landscapes: how cooperation might limit exploration. In *Proc. ACM Genetic and Evolutionary Computation Conference (GECCO'06)*, pages 75–82. ACM.

Wales, D. (2004). *Energy Landscapes : Applications to Clusters, Biomolecules and Glasses (Cambridge Molecular Science)*. Cambridge University Press.

Wales, D. J. and Doye, J. P. K. (1997). Global Optimization by Basin-Hopping and the Lowest Energy Structures of Lennard-Jones Clusters Containing up to 110 Atoms. *J. Phys. Chem. A*, 101(28):5111–5116.