

Alternative Restart Strategies for CMA-ES*

Ilya Loshchilov^{1,2}, Marc Schoenauer^{1,2}, and Michèle Sebag^{2,1}

¹ TAO Project-team, INRIA Saclay - Île-de-France

² Laboratoire de Recherche en Informatique (UMR CNRS 8623)
Université Paris-Sud, 91128 Orsay Cedex, France

FirstName.LastName@inria.fr

Abstract. This paper focuses on the restart strategy of CMA-ES on multi-modal functions. A first alternative strategy proceeds by decreasing the initial step-size of the mutation while doubling the population size at each restart. A second strategy adaptively allocates the computational budget among the restart settings in the BIPOP scheme. Both restart strategies are validated on the BBOB benchmark; their generality is also demonstrated on an independent real-world problem suite related to spacecraft trajectory optimization.

1 Introduction

The long tradition of performance of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) algorithm on real-world problems (with over 100 published applications [6]) is due among others to its good behavior on multi-modal functions. Two versions of CMA-ES with restarts have been proposed to handle multi-modal functions: IPOP-CMA-ES [2] was ranked first on the continuous optimization benchmark at CEC 2005 [4,3]; and BIPOP-CMA-ES [5] showed the best results together with IPOP-CMA-ES on the black-box optimization benchmark (BBOB) in 2009 and 2010.

This paper focuses on analyzing and improving the restart strategy of CMA-ES, viewed as a noisy hyper-parameter optimization problem in a 2D space (population size, initial step-size). Two restart strategies are defined. The first one, NIPOP-aCMA-ES (*New* IPOP-aCMA-ES), differs from IPOP-CMA-ES as it simultaneously increases the population size and decreases the step size. The second one, NBIPOP-aCMA-ES, allocates computational power to different restart settings depending on their current results. While these strategies have been designed with the BBOB benchmarks in mind [8], their generality is shown on a suite of real-world problems [16].

The paper is organized as follows. After describing the weighted active $(\mu/\mu_w, \lambda)$ -CMA-ES and its current restart strategies (section 2), the proposed restart schemes are described in section 3. Section 4 reports on their experimental validation. The paper concludes with a discussion and some perspectives for further research.

* Work partially funded by FUI of System@tic Paris-Region ICT cluster through contract DGT 117 407 *Complex Systems Design Lab* (CSDL).

2 The Weighted Active $(\mu/\mu_w, \lambda)$ -CMA-ES

The CMA-ES algorithm is a stochastic optimizer, searching the continuous space \mathbb{R}^D by sampling λ candidate solutions from a multivariate normal distribution [10,9]. It exploits the best μ solutions out of the λ ones to adaptively estimate the local covariance matrix of the objective function, in order to increase the probability of successful samples in the next iteration. The information about the remaining (worst $\lambda - \mu$) solutions is used only implicitly during the selection process.

In active $(\mu/\mu_I, \lambda)$ -CMA-ES however, it has been shown that the worst solutions can be exploited to reduce the variance of the mutation distribution in unpromising directions [12], yielding a performance gain of a factor 2 for the active $(\mu/\mu_I, \lambda)$ -CMA-ES with no loss of performance on any of tested functions. A recent extension of the $(\mu/\mu_w, \lambda)$ -CMA-ES, *weighted* active CMA-ES [11] (referred to as aCMA-ES for brevity) shows comparable improvements on a set of noiseless and noisy functions from the BOB benchmark suite [7]. In counterpart, aCMA-ES no longer guarantees the covariance matrix to be positive definite, possibly resulting in algorithmic instability. The instability issues can however be numerically controlled during the search; as a matter of fact they are never observed on the BOB benchmark suite.

At iteration t , $(\mu/\mu_w, \lambda)$ -CMA-ES samples λ individuals according to

$$\mathbf{x}_k^{(t+1)} \sim \mathcal{N}\left(\mathbf{m}^{(t)}, \sigma^{(t)2} \mathbf{C}^{(t)}\right), \quad k = 1 \dots \lambda, \tag{1}$$

where $\mathcal{N}(\mathbf{m}, \mathbf{C})$ denotes a normally distributed random vector with mean \mathbf{m} and covariance matrix \mathbf{C} .

These λ individuals are evaluated and ranked, where index $i : \lambda$ denotes the i -th best individual after the objective function. The mean of the distribution is updated and set to the weighted sum of the best μ individuals ($\mathbf{m} = \sum_{i=1}^{\mu} w_i \mathbf{x}_{i:\lambda}^{(t)}$, with $w_i > 0$ for $i = 1 \dots \mu$ and $\sum_{i=1}^{\mu} w_i = 1$).

The active CMA-ES only differs from the original CMA-ES in the adaptation of the covariance matrix $\mathbf{C}^{(t)}$. Like for CMA-ES, the covariance matrix is computed from the best μ solutions, $\mathbf{C}_{\mu}^{+} = \sum_{i=1}^{\mu} w_i \frac{\mathbf{x}_{i:\lambda} - \mathbf{m}^t}{\sigma^t} \times \frac{(\mathbf{x}_{i:\lambda} - \mathbf{m}^t)^T}{\sigma^t}$. The main novelty is to exploit the worst solutions to compute $\mathbf{C}_{\mu}^{-} = \sum_{i=0}^{\mu-1} w_{i+1} \mathbf{y}_{\lambda-i:\lambda} \mathbf{y}_{\lambda-i:\lambda}^T$, where $\mathbf{y}_{\lambda-i:\lambda} = \frac{\|C^{t-1/2}(\mathbf{x}_{\lambda-\mu+1+i:\lambda} - \mathbf{m}^t)\|}{\|C^{t-1/2}(\mathbf{x}_{\lambda-i:\lambda} - \mathbf{m}^t)\|} \times \frac{\mathbf{x}_{\lambda-i:\lambda} - \mathbf{m}^t}{\sigma^t}$. The covariance matrix estimation of these worst solutions is used to decrease the variance of the mutation distribution along these directions:

$$\begin{aligned} \mathbf{C}^{t+1} &= (1 - c_1 - c_{\mu} + c^{-} \alpha_{old}^{-}) \mathbf{C}^t + \\ &+ c_1 \mathbf{p}_c^{t+1} \mathbf{p}_c^{t+1T} + (c_{\mu} + c^{-} (1 - \alpha_{old}^{-})) \mathbf{C}_{\mu}^{+} - c^{-} \mathbf{C}_{\mu}^{-}, \end{aligned} \tag{2}$$

where \mathbf{p}_c^{t+1} is adapted along the evolution path and coefficients c_1 , c_{μ} , c^{-} and α_{old}^{-} are defined such that $c_1 + c_{\mu} - c^{-} \alpha_{old}^{-} \leq 1$. The interested reader is referred to [10,11] for a more detailed description of these algorithms.

As mentioned, CMA-ES has been extended with restart strategies to accommodate multi-modal fitness landscapes, and to specifically handle objective functions with many local optima. As observed by [9], the probability of reaching the optimum (and the overall number of function evaluations needed to do so) is very sensitive to the population size. The default population size $\lambda_{default}$ has been tuned for uni-modal functions; it is hardly large enough for multi-modal functions. Accordingly, [2] proposed a “doubling trick” restart strategy to enforce global search: the restart $(\mu/\mu_w, \lambda)$ -CMA-ES with increasing population, called IPOP-CMA-ES, is a multi-restart strategy where the population size of the run is doubled in each restart until meeting a stopping criterion.

The BIPOP-CMA-ES instead considers two restart regimes. The first one, which corresponds to IPOP-CMA-ES, doubles the population size $\lambda_{large} = 2^{i_{restart}} \lambda_{default}$ in each restart $i_{restart}$ and uses a fixed initial step-size $\sigma_{large}^0 = \sigma_{default}^0$. The second regime uses a small population size λ_{small} and initial step-size σ_{small}^0 , which are randomly drawn in each restart as:

$$\lambda_{small} = \left[\lambda_{default} \left(\frac{1}{2} \frac{\lambda_{large}}{\lambda_{default}} \right)^{U[0,1]^2} \right], \quad \sigma_{small}^0 = \sigma_{default}^0 \times 10^{-2U[0,1]} \quad (3)$$

where $U[0, 1]$ stands for the uniform distribution in $[0, 1]$. Population size λ_{small} thus varies $\in [\lambda_{default}, \lambda_{large}/2]$. BIPOP-CMA-ES launches the first run with default population size and initial step-size. In each restart, it selects the restart regime with less function evaluations. Clearly, the second regime consumes less function evaluations than the doubling regime; it is therefore launched more often.

3 Alternative Restart Strategies

3.1 Preliminary Analysis

The restart strategies of IPOP- and BIPOP-CMA-ES are viewed as a search in the hyper-parameter space.

IPOP-CMA-ES only aims at adjusting population size λ . It is motivated by the results observed on multi-modal problems [9], suggesting that the population size must be sufficiently large to handle problems with global structure. In such cases, a large population size is needed to uncover this global structure and to lead the algorithm to discover the global optimum. IPOP-CMA-ES thus increases the population size in each restart, irrespective of the results observed so far; at each restart, it launches a new CMA-ES with population size $\lambda = \rho_{inc}^{i_{restart}} \lambda_{default}$ (see \circ on Fig. 1). Factor ρ_{inc} must be not too large to avoid “overjumping” some possibly optimal population size λ^* ; it must also be not too small in order to reach λ^* in a reasonable number of restarts. The use of the doubling trick ($\rho_{inc} = 2$) guarantees that the loss in terms of function evaluations (compared to the “oracle“ restart strategy which would directly set the population size to the optimal value λ^*) is about a factor of 2.




Fig. 1. Restart performances in the 2D hyper-parameter space (population size and initial mutation step size in log. coordinates). For each objective function (20 dimensional Rastrigin - top-left, Gallagher 21 peaks - top-right, Katsuuras - bottom-left and Lunacek bi-Rastrigin bottom-right), the median best function value out of 15 runs is indicated. Legends indicate that the optimum up to precision $f(x) = 10^{-10}$ is found always (+), sometimes (\oplus) or never (\circ). Black regions are better than white ones.

On the Rastrigin 20-D function, IPOP-CMA-ES performs well and always finds the optimum after about 5 restarts (Fig. 1, top-left). The Rastrigin function displays indeed a global structure where the optimum is the minimizer of this structure. For such functions, IPOP-CMA-ES certainly is the method of choice. For some other functions such as the Gallagher function, there is no such global structure; increasing the population size does not improve the results. On Katsuuras and Lunacek bi-Rastrigin functions, the optimum can only be found with small initial step-size (lesser than the default one); this explains why it can be solved by BIPOP-CMA-ES, sampling the two-dimensional (λ, σ) space.

Actually, the optimization of a multi-modal function by CMA-ES with restarts can be viewed as the optimization of the function $h(\theta)$, which returns the optimum found by CMA-ES defined by the hyper-parameters $\theta=(\lambda, \sigma)$. Function $h(\theta)$, graphically depicted in Fig. 1 can be viewed as a black box, computationally




Fig. 2. An illustration of λ and σ hyper-parameters distribution for 9 restarts of IPOP-aCMA-ES (\circ), BIPOP-aCMA-ES (\circ and \cdot for 10 runs), NIPOP-aCMA-ES (\square) and NBIPOP-aCMA-ES (\square and many \triangle for $\lambda/\lambda_{default} = 1, \sigma/\sigma_{default} \in [10^{-2}, 10^0]$). The first run of all algorithms corresponds to the point with $\lambda/\lambda_{default} = 1, \sigma/\sigma_{default} = 1$.

expensive and stochastic function (reflecting the stochasticity of CMA-ES). Both IPOP-CMA-ES and BIPOP-CMA-ES are based on implicit assumptions about the $h(\theta)$: IPOP-CMA-ES achieves a deterministic uni-dimensional trajectory, and BIPOP-CMA-ES randomly samples the 2-dimensional search space.

Function $h(\theta)$ also can be viewed as a multi-objective fitness, since in addition to the solution found by CMA-ES, $h(\theta)$ could return the number of function evaluations needed to find that solution. $h(\theta)$ could also return the computational effort SP1 (i.e. the average number of function evaluations of all successful runs, divided by proportion of successful runs). However, SP1 can only be known for benchmark problems where the optimum is known; as the empirical optimum is used in lieu of true optimum, SP1 can only be computed *a posteriori*.

3.2 Algorithm

Two new restart strategies for CMA-ES, respectively referred to as NIPOP-aCMA-ES and NBIPOP-aCMA-ES, are presented in this paper.

If the restart strategy is restricted to the case of increasing of population size (IPOP), we propose to use NIPOP-aCMA-ES, where we additionally decrease the initial step-size by some factor $\rho_{\sigma dec}$. The rationale behind this approach is that the CMA-ES with relatively small initial step-size is able to explore small basins of attraction (see Katsuuras and Lunacek bi-Rastrigin functions on Fig. 1), while with initially large step-size and population size it will neglect the local structure of the function, but converge to the minimizer of the global structure. Moreover, initially, relatively small step-size will quickly increase if it makes sense, and this will allow the algorithm to recover the same global search properties than with initially large step-size (see Rastrigin function on Fig. 1).

NIPOP-CMA-ES thus explores the two-dimensional hyper-parameter space in a deterministic way (see \square symbols on Fig. 2). For $\rho_{\sigma dec} = 1.6$ used in this

study, NIPOP-CMA-ES thus reaches the lower bound ($\sigma = 10^{-2}\sigma_{default}$) used by BIPOP-CMA-ES after 9 restarts, expectedly reaching the same performance as BIPOP-CMA-ES albeit it uses only a large population.

The second restart strategy, NBIPOP-aCMA-ES, addresses the case where the probability to find the global optimum does not much vary in the (λ, σ) space. Under this assumption, it makes sense to have many restarts for a fixed budget (number of function evaluations). Specifically, NBIPOP-aCMA-ES implements the competition of the NIPOP-aCMA-ES strategy (increasing λ and decreasing initial σ^0 in each restart) and a uniform sampling of the σ space, where λ is set to $\lambda_{default}$ and $\sigma^0 = \sigma_{default}^0 \times 10^{-2U[0,1]}$. The selection between the two (NIPOP-aCMA-ES and the uniform sampling) depends on the allowed budget like in NBIPOP-aCMA-ES. The difference is that NBIPOP-aCMA-ES adaptively sets the budget allowed to each restart strategy, where the restart strategy leading to the overall best solution found so far is allowed twice ($\rho_{budget} = 2$) a budget compared to the other strategy.

4 Experimental Validation

The experimental validation of NIPOP-aCMA-ES and NBIPOP-aCMA-ES investigates the performance of the approach comparatively to IPOP-aCMA-ES and BIPOP-aCMA-ES on BBOB noiseless problems and one black-box real-world problem related to spacecraft trajectory optimization. The default parameters of CMA-ES [11,5] are used. This section also presents the first experimental study of BIPOP-aCMA-ES¹, the active version of BIPOP-CMA-ES [5].

4.1 Benchmarking with BBOB Framework

The BBOB framework [7] is made of 24 noiseless and 30 noisy functions [8]. Only the noiseless case has been considered here. Furthermore, only the 12 multimodal functions among these 24 noiseless functions are of interest for this study, as CMA-ES can solve the 12 other functions without any restart.

With same experimental methodology as in [7], the results obtained on these benchmark functions are presented in Fig. 4 and Table 1. The results are given for dimension 40, because the differences are larger in higher dimensions. The **expected running time (ERT)**, used in the figures and table, depends on a given target function value, $f_t = f_{opt} + \Delta f$. It is computed over all relevant trials as the number of function evaluations required in order to reach f_t , summed over all 15 trials, and divided by the number of trials that actually reached f_t [7].

NIPOP-aCMA-ES. On 6 out of 12 test functions ($f_{15}, f_{16}, f_{17}, f_{18}, f_{23}, f_{24}$) NIPOP-aCMA-ES obtains the best known results for BBOB-2009 and BBOB-2010 workshops. On f_{23} Katsuuras and f_{24} Lunacek bi-Rastrigin, NIPOP-aCMA-ES has a speedup of a factor from 2 to 3, as could have been expected. It performs

¹ For the sake of reproducibility, the source code for NIPOP-aCMA-ES and NBIPOP-aCMA-ES is available at <https://sites.google.com/site/ppsnbipop/>

unexpectedly well on f_{16} Weierstrass functions, 7 times faster than IPOPOP-aCMA-ES and almost 3 times faster than BIPOP-aCMA-ES. Overall, according to Fig. 4, NIPOP-aCMA-ES performs as well as BIPOP-aCMA-ES, while restricted to only one regime of increasing population size.

NBIPOP-aCMA-ES. Thanks to the first regime of increasing population size, NBIPOP-aCMA-ES inherits some results of NIPOP-aCMA-ES. However, on functions where the population size does not play any important role, it performs significantly better than BIPOP-aCMA-ES. This is the case for f_{21} Gallagher 101 peaks and f_{22} Gallagher 21 peaks functions, where NBIPOP-aCMA-ES has a speedup of a factor of 6. It seems that the adaptive choice between two regimes works efficiently on all functions except on f_{16} Weierstrass. In this last case, NBIPOP-aCMA-ES mistakingly prefers small populations, with a loss factor 4 compared to NIPOP-aCMA-ES. According to Fig. 4, NBIPOP-aCMA-ES performs better than BIPOP-aCMA-ES on weakly structured multi-modal functions, showing overall best results for BBOB-2009 and BBOB-2010 workshops in dimensions 20 (results not shown here) and 40.

Due to space limitations, the interested reader is referred to [13] for a detailed presentation of the results.

4.2 Interplanetary Trajectory Optimization

The NIPOP-aCMA-ES and NBIPOP-aCMA-ES strategies, designed for the BBOB benchmark functions, can possibly *overfit* this benchmark suite. In order to test the generality of these strategies, a real-world black-box problem is considered, pertaining to a completely different domain: Advanced Concepts Team of European Space Agency is making available several difficult spacecraft trajectory optimization problems as black box functions to invite the operational research community to compare different derivative-free solvers on these test problems [16].

The following results consider the 18-dimensional bound-constrained black-box function “TandEM-Atlas501”, that defines an interplanetary trajectory to Saturn from the Earth with multiple fly-bys, launched by the rocket Atlas 501. The final goal is to maximize the mass $f(x)$, which can be delivered to Saturn using one of 24 possible fly-by sequences with possible maneuvers around Venus, Mars and Jupiter.

The first best results was found for a sequence Earth-Venus-Earth-Earth-Saturn ($f_{max} = 1533.45$) in 2008 by B. Addis et al. [1]. The best results so far ($f_{max} = 1673.88$) was found in 2011 by G. Stracquadanio et al. [15].

All versions of CMA-ES with restarts have been launched with a maximum budget of 10^8 function evaluations. All variables are normalized in the range $[0, 1]$. In the case of sampling outside of boundaries, the fitness is penalized and becomes $f(x) = f(x_{feasible}) - \alpha \|x - x_{feasible}\|^2$, where $x_{feasible}$ is the closest feasible point from point x and α is a penalty factor, which was arbitrarily set to 1000.

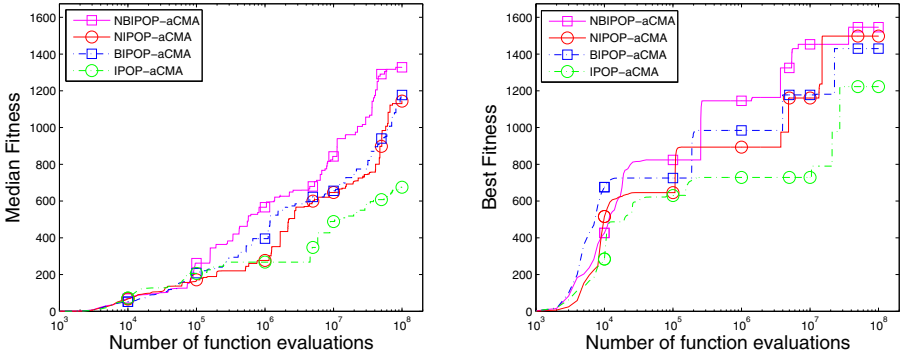


Fig. 3. Comparison of all CMA-ES restart strategies on the Tandem fitness function (mass): median (left) and best (right) values out of 30 runs

As shown on Fig. 3, the new restart strategies NIPOP-aCMA-ES and NBIPOP-aCMA-ES respectively improve on the former ones (IPOP-aCMA-ES and BIPOP-aCMA-ES); further, NIPOP-aCMA-ES reaches same performances as BIPOP-aCMA-ES.

The best solution found by NBIPOP-aCMA-ES² improves on the best solution found in 2008, while it is worse than the current best solution, which is blamed on the lack of problem specific heuristics [1,15], on the possibly insufficient time budget (10^8 fitness evaluations), and also on the lack of appropriate constraint handling heuristics.

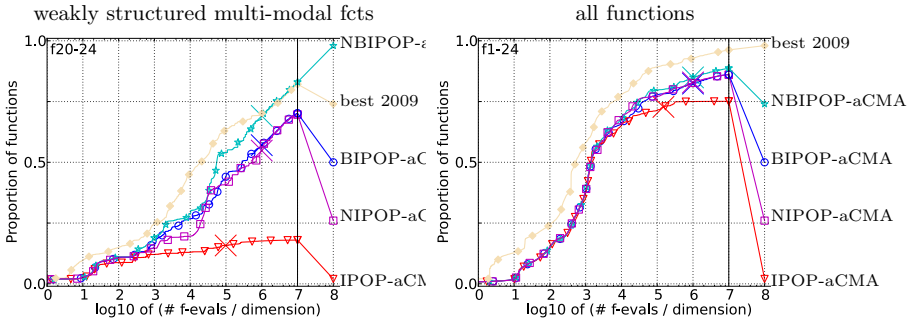


Fig. 4. Bootstrapped empirical cumulative distribution of the number of objective function evaluations divided by dimension (FEvals/D) for 50 targets in $10^{[-8..2]}$ for all functions and weakly structured multi-modal subgroup in 40-D. The “best 2009” line corresponds to the best ERT observed during BBOB 2009 for each single target.

² $x = [0.83521, 0.45092, 0.50284, 0.65291, 0.61389, 0.75773, 0.43376, 1, 0.89512, 0.77264, 0.11229, 0.20774, 0.018255, 6.2057e-09, 4.0371e-08, 0.2028, 0.36272, 0.32442]$; $fitness(x) = mass(x) = 1546.5$.

Table 1. Overall results on multi-modal functions $f3 - 4$ and $f15 - 24$ in dimension $d = 40$: Expected running time (ERT in number of function evaluations) divided by the respective best ERT measured during BBOB-2009 for precision Δf ranging in 10^i , $i = 1 \dots -7$. The median number of conducted function evaluations is additionally given in *italics*, if $\text{ERT}(10^{-7}) = \infty$. #succ is the number of trials that reached the final target $f_{\text{opt}} + 10^{-8}$. Best results are printed in bold. For a more detailed (statistical) analysis of results on BBOB problems, please see [13]. Statistically significantly better entries (Wilcoxon rank-sum test with $p = 0.05$) are indicated in bold. The interested reader is referred to [13] for the statistical analysis and discussion of these results.

Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f3	15526	15602	15612	15646	15651	15656	15/15
BIPOP-a	2395	∞	∞	∞	∞	∞	4e7 / 0/15
IPOP-aC	∞	∞	∞	∞	∞	∞	6e6 / 0/8
NBIPOP-a	8177	∞	∞	∞	∞	∞	4e7 / 0/15
NIPOP-a	4615	∞	∞	∞	∞	∞	4e7 / 0/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f4	15536	15601	15659	15703	15733	2.8e5	6/15
BIPOP-a	∞	∞	∞	∞	∞	∞	4e7 / 0/15
IPOP-aC	∞	∞	∞	∞	∞	∞	6e6 / 0/8
NBIPOP-a	∞	∞	∞	∞	∞	∞	4e7 / 0/15
NIPOP-a	∞	∞	∞	∞	∞	∞	4e7 / 0/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f15	1.9e5	7.9e5	1.0e6	1.1e6	1.1e6	1.1e6	15/15
BIPOP-a	1.2	1.1	1.1	1.1	1.1	1.1	15/15
IPOP-aC	0.72	0.43	0.60	0.61	0.62	0.63	8/8
NBIPOP-a	1.0	0.71	0.75	0.76	0.77	0.77	15/15
NIPOP-a	0.92	0.61	0.55	0.56	0.57	0.58	15/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f16	5244	72122	3.2e5	1.4e6	2.0e6	2.0e6	15/15
BIPOP-a	1.3	0.96	0.80	0.54	0.50	0.51	15/15
IPOP-aC	0.91	1.1	1.0	0.51	1.4	1.4	8/8
NBIPOP-a	0.97	0.78	0.34	0.38	0.46	0.74	15/15
NIPOP-a	1.2	0.65	0.23	0.21	0.16	0.18	15/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f17	399	4220	14158	51958	1.3e5	2.7e5	14/15
BIPOP-a	1.1	0.64	1.6	1.1	1.4	0.87	15/15
IPOP-aC	1.0	0.52	1.3	1.3	0.97	0.83	8/8
NBIPOP-a	1.0	0.57	1.2	1.2	1.0	0.81	15/15
NIPOP-a	0.97	0.52	0.97	1.00	1.1	0.70	15/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f18	1442	16998	47068	1.9e5	6.7e5	9.5e5	15/15
BIPOP-a	0.94	0.51	1.0	0.98	0.88	0.67	15/15
IPOP-aC	0.96	0.68	1.0	0.66	0.45	0.48	8/8
NBIPOP-a	1.0	0.97	1.1	0.93	0.57	0.53	15/15
NIPOP-a	0.95	0.58	0.75	0.71	0.50	0.42	15/15

Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f19	1	1	1.4e6	2.6e7	4.5e7	4.5e7	8/15
BIPOP-a	396	6.7e4	0.87	1.2	1.0	1.0	9/15
IPOP-aC	462	4.4e40.57	0.34	0.20	0.20	0.20	8/8
NBIPOP-a	424	8.3e4	0.97	0.81	1.1	1.1	9/15
NIPOP-a	436	8.2e4	1.9	0.48	0.32	0.32	15/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f20	222	1.3e5	1.6e8	∞	∞	∞	0
BIPOP-a	4.0	9.0	0.34	.	.	.	0/15
IPOP-aC	3.9	8.1	0.18	.	.	.	0/8
NBIPOP-a	4.0	8.5	0.39	.	.	.	0/15
NIPOP-a	4.0	6.5	0.32	.	.	.	0/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f21	1044	21144	1.0e5	1.0e5	1.0e5	1.0e5	26/30
BIPOP-a	7.5	60	37	37	37	37	15/15
IPOP-aC	7.1	421	∞	∞	∞	∞	3e6 / 0/8
NBIPOP-a	4.9	10	5.1	5.1	5.1	5.1	15/15
NIPOP-a	14	440	173	172	171	171	12/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f22	3090	35442	6.5e5	6.5e5	6.5e5	6.5e5	8/30
BIPOP-a	12	343	201	200	200	199	4/15
IPOP-aC	144	93	∞	∞	∞	∞	3e6 / 0/8
NBIPOP-a	12	112	32	32	32	32	12/15
NIPOP-a	179	583	∞	∞	∞	∞	4e7 / 0/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f23	7.1	11925	75453	1.3e6	3.2e6	3.4e6	15/15
BIPOP-a	8.4	7.8	1.3	1.9	1.00	0.99	15/15
IPOP-aC	9.2	∞	∞	∞	∞	∞	4e6 / 0/8
NBIPOP-a	8.6	10	1.6	1.3	0.58	0.59	15/15
NIPOP-a	5.9	61	11	0.72	0.36	0.38	15/15
Δf_{opt}	1e1	1e0	1e-1	1e-3	1e-5	1e-7	#succ
f24	5.8e6	9.8e7	3.0e8	3.0e8	3.0e8	3.0e8	1/15
BIPOP-a	3.6	1.4	∞	∞	∞	∞	4e7 / 0/15
IPOP-aC	∞	∞	∞	∞	∞	∞	1e7 / 0/8
NBIPOP-a	2.1	0.19	0.97	0.97	0.97	0.97	2/15
NIPOP-a	1.2	0.15	0.44	0.44	0.44	0.44	4/15

5 Conclusion and Perspectives

This paper contribution regards two new restart strategies for CMA-ES. NIPOP-aCMA-ES is a deterministic strategy simultaneously increasing the population size *and* decreasing the initial step-size of the Gaussian mutation. NBIPOP-aCMA-ES implements a competition between NIPOP-aCMA-ES and a random sampling of the initial mutation step-size, adaptively adjusting the computational budget of each one depending on their current best results. Besides the extensive validation of NIPOP-aCMA-ES and NBIPOP-aCMA-ES on the BBOB benchmark, the generality of these strategies has been tested on a new problem, related to interplanetary spacecraft trajectory planning.

The main limitation of the proposed restart strategies is to quasi implement a deterministic trajectory in the θ space. Further work will consider $h(\theta)$ as yet another expensive noisy black-box function, and the use of a CMA-ES in the hyper-parameter space will be studied. The critical issue is naturally to keep

the overall number of fitness evaluations beyond reasonable limits. A surrogate-based approach will be investigated [14], learning and exploiting an estimate of the (noisy and stochastic) $h(\theta)$ function.

References

1. Addis, B., Cassioli, A., Locatelli, M., Schoen, F.: Global optimization for the design of space trajectories. *Optimization On Line*, 11 (2008)
2. Auger, A., Hansen, N.: A restart CMA evolution strategy with increasing population size. In: *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2005)*, pp. 1769–1776. IEEE Press (2005)
3. García, S., Molina, D., Lozano, M., Herrera, F.: A study on the use of non-parametric tests for analyzing the evolutionary algorithms' behaviour: a case study on the CEC 2005 special session on real parameter optimization. *Journal of Heuristics* 15, 617–644 (2009)
4. Hansen, N.: *Compilation of results on the 2005 CEC benchmark function set* (May 2006)
5. Hansen, N.: Benchmarking a BI-population CMA-ES on the BBOB-2009 function testbed. In: Rothlauf, F. (ed.) *GECCO Companion*, pp. 2389–2396. ACM (2009)
6. Hansen, N.: *References to CMA-ES applications* (2009), <http://www.lri.fr/~hansen/cmaapplications.pdf>
7. Hansen, N., Auger, A., Finck, S., Ros, R.: *Real-parameter black-box optimization benchmarking 2012: Experimental setup*. Technical report, INRIA (2012)
8. Hansen, N., Finck, S., Ros, R., Auger, A.: *Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions*. Technical Report RR-6829, INRIA (2009) (updated, February 2010)
9. Hansen, N., Kern, S.: *Evaluating the CMA Evolution Strategy on Multimodal Test Functions*. In: Yao, X., Burke, E.K., Lozano, J.A., Smith, J., Merelo-Guervós, J.J., Bullinaria, J.A., Rowe, J.E., Tiño, P., Kabán, A., Schwefel, H.-P. (eds.) *PPSN 2004*. LNCS, vol. 3242, pp. 282–291. Springer, Heidelberg (2004)
10. Hansen, N., Ostermeier, A.: *Completely derandomized self-adaptation in evolution strategies*. *Evolutionary Computation* 9(2), 159–195 (2001)
11. Hansen, N., Ros, R.: *Benchmarking a weighted negative covariance matrix update on the BBOB-2010 noiseless testbed*. In: *GECCO 2010: Proceedings of the 12th Annual Conference Comp. on Genetic and Evolutionary Computation*, pp. 1673–1680. ACM, New York (2010)
12. Jastrebski, G.A., Arnold, D.V.: *Improving evolution strategies through active covariance matrix adaptation*. In: *IEEE Congress on Evolutionary Computation, CEC 2006*, pp. 2814–2821 (2006)
13. Loshchilov, I., Schoenauer, M., Sebag, M.: *Black-box Optimization Benchmarking of NIPOP-aCMA-ES and NBIPOP-aCMA-ES on the BBOB-2012 Noiseless Testbed*. In: *GECCO 2012: Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*. ACM (page to appear, 2012)
14. Loshchilov, I., Schoenauer, M., Sebag, M.: *Self-Adaptive Surrogate-Assisted Covariance Matrix Adaptation Evolution Strategy*. In: *GECCO 2012 Proceedings*. ACM (page to appear, 2012)
15. Stracquadanio, G., La Ferla, A., De Felice, M., Nicosia, G.: *Design of robust space trajectories*. In: *Research and Development in Intelligent Systems XXVIII*, pp. 341–354. Springer (2011)
16. Vinko, T., Izzo, D.: *Global Optimisation Heuristics and Test Problems for Preliminary Spacecraft Trajectory Design*. Technical Report GOHTPPSTD, European Space Agency (2008)