

A Continuous Variable Neighbourhood Search Based on *Specialised EAs*: Application to the Noiseless BBO-Benchmark 2009

Carlos García-Martínez
Dept. of Computing and Numerical Analysis
14071 University of Córdoba
Córdoba, Spain
gcarcia@uco.es

Manuel Lozano
Dept. of Computer Science and Artificial
Intelligence
CITIC-UGR (Research Center on Information
and Communications Technology)
18071 University of Granada
Granada, Spain
lozano@decsai.ugr.es

ABSTRACT

Variable Neighbourhood Search is a metaheuristic combining three components: generation, improvement, and shaking components. In this paper, we design a continuous Variable Neighbourhood Search algorithm based on three *specialised Evolutionary Algorithms*, which play the role of each aforementioned component: 1) an EA specialised in generating a good starting point as generation component, 2) an EA specialised in exploiting local information as improvement component, 3) and another EA specialised in providing local diversity as shaking component. Experiments are carried out on the noiseless Black-Box Optimisation Benchmark 2009 testbed.

Categories and Subject Descriptors

G.1.6 [Numerical Analysis]: Optimization Global Optimization, Unconstrained Optimization; F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems

General Terms

Algorithms

Keywords

Benchmarking, Black-box optimization, Evolutionary computation, Hybrid metaheuristics, Variable neighbourhood search, Specialised evolutionary algorithms

1. INTRODUCTION

Metaheuristics (MHs) [12] are a family of search and optimisation algorithms based on extending basic heuristic meth-

ods by including them into an iterative framework augmenting their exploration capabilities. MHs coordinate *subordinate components* (such as probability distributions, tabu lists or genetic operators among others) with the aim of performing an effective and efficient process in searching for the global optimum of a problem.

Over the last years, a large number of search algorithms were reported that do not purely follow the concepts of one single classical MH, but they attempt to obtain the best from a set of MHs that perform together and complement each other to produce a profitable synergy from their combination. These approaches are commonly referred to as *hybrid MHs* [26].

Evolutionary Algorithms (EAs) [5] are stochastic search methods that mimic the metaphor of natural biological evolution. EAs rely on the concept of a population of individuals, which undergo probabilistic operators to evolve toward increasingly better fitness values of the individuals.

A novel method to build hybrid MHs, recently introduced in [23], concerns the incorporation of *specialised EAs* into classical MHs, replacing determinate components, but preserving the essence of the original MH as much as possible. The idea is to build customised EAs playing the same role as particular MH components, but more effectively, i.e., *evolutionary MH components*. In this way, a classical MH is transformed into an integrative hybrid MH (because one of its components is another MH). In the literature, we find some proposals that follow this idea: In [1] and [9], two EAs are applied to perform local search processes within a multi-start MH; and in [22, 23], an evolutionary ILS-perturbation technique is presented (ILS is for Iterated Local Search).

Variable Neighbourhood Search (VNS) [25] is a MH that exploits systematically the idea of neighbourhood change (from a given set of neighbourhood structures N_k , $k = 1, \dots, k_{max}$), both in the descent to local minima and in the escape from the valleys which contain them. It mainly consists of the following three components, which work on a single candidate solution, the *current solution* (s^c) (see Fig. 1):

1. *Generation component*: Firstly, a method is executed to generate s^c within the search space.
2. *Improvement component*: Secondly, s^c is refined, usually by a local search method.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'09, July 8–12, 2009, Montréal Québec, Canada.
Copyright 2009 ACM 978-1-60558-505-5/09/07 ...\$5.00.

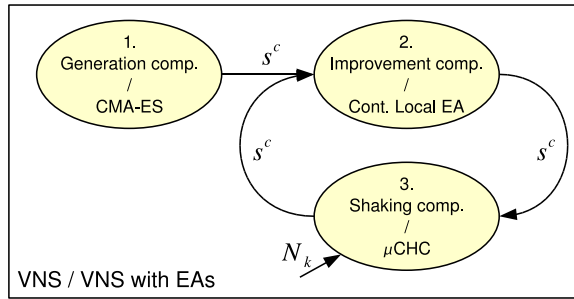


Figure 1: General VNS model / VNS based on specialised EAs

3. *Shaking component*: Then, shaking is performed to escape from the valley where s^c lies. It selects a random solution in the k th neighbourhood of s^c , which becomes a new starting solution for improvement component. At the beginning of the run, and every time last improvement process improved the best found solution, k is set to one; otherwise, k is set to $k + 1$.

In this study, we design a continuous VNS model based on three specialised EAs, which play the role of each VNS component. The algorithm is benchmarked on the noiseless Black-Box Optimization Benchmark 2009 testbed [13].

In Sect. 2, we describe the proposed VNS model. In Sect. 3, we present the experimental results. In Sect. 4, CPU timing experiment is shown. Conclusions are presented in Sect. 5.

2. VNS BASED ON SPECIALISED EAS

In this study, we design a continuous VNS model based on three specialised EAs, which play the role of each VNS component (see Fig. 1):

1. *Covariance Matrix Adaptation Evolution Strategy* (CMA-ES) [16, 17] generates initial s^c (Sect. 2.1).
2. *Continuous Local EA* improves s^c (Sect. 2.2).
3. *Micro Cross-generational elitist sel., Heterogeneous recombination, and Cataclysmic mutation* (μ CHC) performs shaking to scape from the valley where s^c lies, according to k th neighbourhood (Sect. 2.3).

Adaptation of parameter k is performed the same way general VNS model does. However, in order to avoid frequent, but too small improvements when tackling continuous optimisation problems, our algorithm will set k to one only if the improvement is superior to a given threshold ($1e - 8$). Otherwise, k is set to $k + 1$.

Neighbourhoods structures are defined by using metric ρ :

$$N_k(s^c) = \{s \mid r_{k-1} \leq \rho(s^c, s) \leq r_k\} \quad (1)$$

where r_k is the radius of $N_k(s^c)$ monotonically nondecreasing with k . Hereafter, we will always consider the following distance metric because of computational reasons:

$$\rho(x, y) = \sum_{i=1}^n |x_i - y_i| \quad (2)$$




Figure 2: Continuous Local EA

Stop condition is reached when k surpasses k_{max} (in our case, k_{max} is set to 20). At this point, as it is suggested in [13], the algorithm performs a *restart*. In this way, exploiting larger number of function evaluations may increase the chance to achieve better function values or solve the function up to the final target. In our case, a restart means that VNS is run once again from the CMA-ES execution, generating a new s^c .

2.1 CMA-ES as Generator Component

CMA-ES [16, 17] was originally introduced to improve the local performances of evolution strategies, however, it even reveals competitive global search performances [15]. In CMA-ES, not only is the step size of the mutation operator adjusted at each generation, but so too is the step direction in the multidimensional problem space, i.e., not only is there a mutation strength per dimension but their combined update is controlled by a covariance matrix whose elements are updated as the search proceeds.

In this work, we apply CMA-ES to obtain a good point in the design space from which the algorithm carries out its search process. We have applied the C code, available at http://www.lri.fr/~hansen/maes_inmatlab.html with the suggested parameter values from a random solution.

2.2 Continuous Local EA as Improvement Component

In this work, we propose to apply *Continuous Local EA* as the improvement component of VNS. It is an EA based on the principles of Binary Local Genetic Algorithm [11], which obtained promising results in efficacy and efficiency against classical improvement methods.

Continuous Local EA is a steady-state *real-coded genetic algorithm* [18] that inserts one single new member into the population (Pop) in each iteration. It uses a replacement method in order to force a member of the current population to perish and to make room for the new offspring. It is important to know that the selected replacement method favours the acquisition of information about the search space. Then, the algorithm exploits local information around s^c , provided by the individuals in Pop close to it, to orientate the improvement process.

Let's suppose that Continuous Local EA is to improve the given s^c . Then, following steps are carried out (see Fig. 2):

- *Mate selection*: A solution from Pop is selected as s^{mate} by means of *nearest improving solution selection method* (Sect. 2.2.1).
- *Crossover*: s^{mate} and s^c are crossed over by means of *parent-centric BLX- α* , creating offspring z (Sec. 2.2.2).

- *Acceptance and replacement*: If z is accepted (Sect. 2.2.3), it replaces s^c and is inserted into the population forcing the individual with lowest information contribution to perish (Sect. 2.2.4).

All these steps are repeated until a maximum number of iterations without improving the best visited s^c is reached (we stop the run after 100 iterations without improving best s^c). At last, the best found s^c is returned.

An important aspect when using Continuous Local EA in the proposed VNS model is that *Pop* undergoes initialisation only once, at the beginning of the run, and not at every invocation to improve s^c . In this way, the algorithm may gather up valuable information about the search space and its optima, and employ accumulated search experience from previous refinements to enhance future ones.

2.2.1 Nearest Improving Solution Selection

Assortative mating is the natural occurrence of mating between individuals of similar phenotype more or less often than expected by chance. Mating between individuals with similar phenotype more often is called positive assortative mating and less often is called negative assortative mating. Fernandes et al. [7] implement these ideas to design two mating selection mechanisms.

We introduce a new assortative mating mechanism to be used in Continuous Local EA, *nearest improving solution selection*. Its idea is to select the individual in *Pop* that provides valuable information of the search space (location and fitness value) for the task of locally improving s^c . On the one hand, this mechanism dismisses individuals worse than s^c , because they reduce chances for improving s^c . On the other hand, nearest solution is selected because it provides more information about the search space region where s^c is located. We also include a distance threshold to avoid choosing a solution too similar to s^c .

Therefore, first parent is always s^c , then, second parent is the individual in *Pop* most similar to s^c (at phenotypic level) fulfilling two conditions: 1) it has a better fitness value than s^c and 2) similarity between both solutions is above a given threshold. Mating threshold has been initially set to $1e - 2$. Then, after every VNS restart, it is set to half its value to perform a more precise search.

An important remark is that this selection mechanism may return no mate for s^c . It occurs when the mechanism has dismissed all solutions in *Pop* because they are either worse than s^c , either their distances to s^c are under mating threshold. Crossover is not performed in this case. Offspring is generated by perturbing s^c instead (Sect. 2.2.2).

2.2.2 PBX- α Crossover Operator

PBX- α [24] is an instance of *parent-centric crossover operators*, which have arisen as a meaningful and efficient way of solving real-parameter optimisation problems [4, 10].

Given $s^c = (s_1^c \dots s_n^c)$ and $s^{mate} = (s_1^{mate} \dots s_n^{mate})$, PBX- α generates offspring $z = (z_1 \dots z_n)$, where z_i is a random (uniform) number from interval $[s_i^c - I \cdot \alpha, s_i^c + I \cdot \alpha]$, with $I = |s_i^c - s_i^{mate}|$. Parameter α is initially set to 0.5. Then, after every VNS restart, it is set to $0.5/\log_e(\text{number of restarts} + 1)$ for a more precise search.

As mentioned before, crossover operation occurs when mate selection returns a solution from *Pop*. Otherwise, offspring is generated by adding a normal random vector to s^c ,

whose standard deviation is the product of the current values of aforementioned parameters α and mating threshold.

2.2.3 Acceptance Criterion

Once offspring z has been generated either by PBX- α or perturbation, Continuous Local EA decides which solution, between z and s^c , becomes the new s^c . It follows a similar idea to the ones behind *Simulated Annealing* [20] and *Noising Methods* [2] to overcome rugged landscapes and local optima. In particular, it always accepts z if it is better than a specific computed bound. At the beginning, that bound is set equal to the fitness value of s^c , only allowing improving offspring. Then, every time a new best s^c is found, the bound is computed as the average of the fitness of the new best s^c and the old bound. This simple mechanism provides some selection pressure on the search process that also lets the algorithm to accept uphill moves, and overcome small local optima.

2.2.4 Lowest Information Contribution Replacement

One of the objectives of Continuous Local EA is to gather up valuable information about the search space and its optima, to enhance future refinements. To carry out this objective, *Pop* should maintain a set of good solutions properly spread over the search space. The problem of maintaining such a set of solutions recalls the problem of attaining a good set of non-dominated solutions in multi-objective problems [3]. In [21], Kukkonen and Deb propose a method for pruning of non-dominated solutions in many-objective problems. The basic idea is to eliminate the most crowded members of a non-dominated set one by one, and update the crowding information of the remaining members after each removal.

In this work, we propose a similar method for updating *Pop* of Continuous Local EA, when solving mono-objective problems. The mechanism firstly inserts the new offspring in *Pop*, then, it removes the most crowded solution, keeping population size constant (we use $|Pop| = 100$). To determine the solution to be removed, everyone in *Pop* holds a pointer to its nearest solution as well as the distance between them, at phenotypic level. Since, according to this definition, there are always, at least, two solutions holding the minimum distance, the worst of them is the one to be removed. Pointers and distances are updated every time *Pop* changes, i.e., when inserting the new solution and when removing the most crowded one. Notice that the best found solution is never removed.

2.3 μ CHC as Shaking Component

In this work, we propose μ CHC as the shaking component of VNS. μ CHC was firstly presented as an *evolutionary ILS-perturbation technique* for binary-coded problems [22, 23]. The role of μ CHC is to receive s^c , provide local diversity, and generate a solution that is then considered as starting point for the next improvement process (see Fig. 3). The reason for choosing CHC, as the base model, is that it suitably combines powerful diversification mechanisms with an elitist selection strategy. The filtering of high diversity by means of high selective pressure favours the creation of *useful diversity*: many dissimilar solutions are produced during the run and only the best ones are conserved in the population, allowing diverse and promising solutions to be maintained. From our point of view, this behaviour is desirable for an EA assuming the work of a shaking operator.




Figure 3: μ CHC as shaking module

2.3.1 μ CHC Model

We have conceived μ CHC to be an effective *explorer* in the neighbourhood of s^c to perform shaking process, because it provides local diversity. At the beginning, s^c is used to create its initial population (Sect. 2.3.2). Then, it is performed throughout a predetermined number of fitness function evaluations. The best reached individual is then considered as starting point for the next improvement process (see Fig. 3).

The main components of the algorithm are:

- *Population size:* μ CHC manages a population with few individuals ($|Pop| = 5$), and thus, it may be seen as micro EA. In standard VNS models, the number of evaluations required by the shaking mechanism is very low as compared with the one for the improvement method. With the aim of preserving, as far as possible, the essence of VNS, we have considered an EA with a low sized population; for being able to work adequately under the requirement of spending few evaluations.
- *Number of evaluations:* In particular, the number of evaluations assigned to μ CHC for a particular invocation will be a fixed proportion (p_{evals}), of the number of evaluations consumed by the previously performed improvement method (Continuous Local EA). It is worth noting that p_{evals} should be set to a low value. We will use p_{evals} equals to 0.5.
- *Elitist selection:* Current population is merged with offspring population obtained from it and the best $|Pop|$ individuals are selected for the new population.
- *Incest prevention mechanism:* Before mating, Hamming distance between the corresponding Gray-coding strings of paired individuals (with 20 bits per variable) is calculated. Only paired individuals whose distance exceed a difference threshold d are allowed to undergo crossover operation. Aforementioned threshold is initialised to $L/4$ (with L being the number of bits coding a potential solution). If no offspring is obtained in one generation, difference threshold is decremented by one.
- *BLX- α crossover operator:* Paired individuals allowed to produce offspring undergo BLX- α crossover operator [6], producing one offspring. Given y^1 and y^2 , BLX- α generates offspring z , where z_i is a random (uniform) number from $[y^{min} - I \cdot \alpha, y^{max} + I \cdot \alpha]$ interval, with $y^{min} = \min(y_i^1, y_i^2)$, $y^{max} = \max(y_i^1, y_i^2)$, and $I = |y_i^1 - y_i^2|$. This parameter α is set to 0.5.
- *Mating with s^c :* μ CHC incorporates the strategy of recombining s^c with another solution [19]. In addition to




Figure 4: Shape of probability distribution for initialisation

the typical recombination phase of CHC, our algorithm mates s^c with a member of Pop (selected at random) and, if they are finally crossed over (attending on the incest prevention mechanism), the resulting offspring will be introduced into the offspring population.

- *Cataclysmic mutation:* μ CHC, as CHC, uses no mutation in the classical sense of the concept, but instead, it goes through a process of cataclysmic mutation when the population has converged. The difference threshold is considered to measure the stagnation of the search, which happens when it has dropped to one. Then, the population is reinitialised with individuals generated by perturbing s^c attending to the k th neighbourhood structure (Sect. 2.3.2).

2.3.2 Initialisation and Cataclysmic Mutation

Every individual in the initial μ CHC population is generated by perturbing s^c according to a relaxed idea of the given neighbourhood structure N_k (see equation 1). In particular, each individual is generated by adding a vector of random values from the following equation (see Fig. 4):

$$N(0, 1) \cdot range_i / k_{max} \pm k \cdot range_i / k_{max} \quad (3)$$

where $N(0, 1)$ is a normal random value with mean 0 and variance 1, $range_i$ is the range of i th variable domain, k_{max} is the maximum number of neighbourhood structures VNS manages, and k is the current neighbourhood index; \pm is chosen according to the sign of the random value.

Cataclysmic mutation fills the population with individuals created by the same way as initial population is built, but preserving the best performing individual found in the previous generation. After applying cataclysmic mutation, the difference threshold is set to: $\sigma \cdot (1 - \sigma) \cdot L$, with L being the number of bits coding a potential solution of the problem (20 bits per variable) and σ is the average of the minimal radius of neighbourhoods N_k and N_{k+1} ($(2k + 1) / (2k_{max})$).

3. RESULTS

Results from experiments according to [13] on the benchmark functions given in [8, 14] are presented in Figures 5 and 6 and in Table 1.

Individual experiments were run with a maximum and adaptive time budget equal to the available time (from 2009-04-07 to 2009-04-16) divided by the number of remaining experiments (initially, $5dims \cdot 24funcs \cdot 15runs$). *Crafting effort* [13] was 0 because all the functions where approached with the very same parameter setting.

4. CPU TIMING EXPERIMENT

The continuous Variable Neighbourhood Search based on specialised Evolutionary Algorithms was run with a maximum of $10^5 \cdot D$ function evaluations and restarted until 30

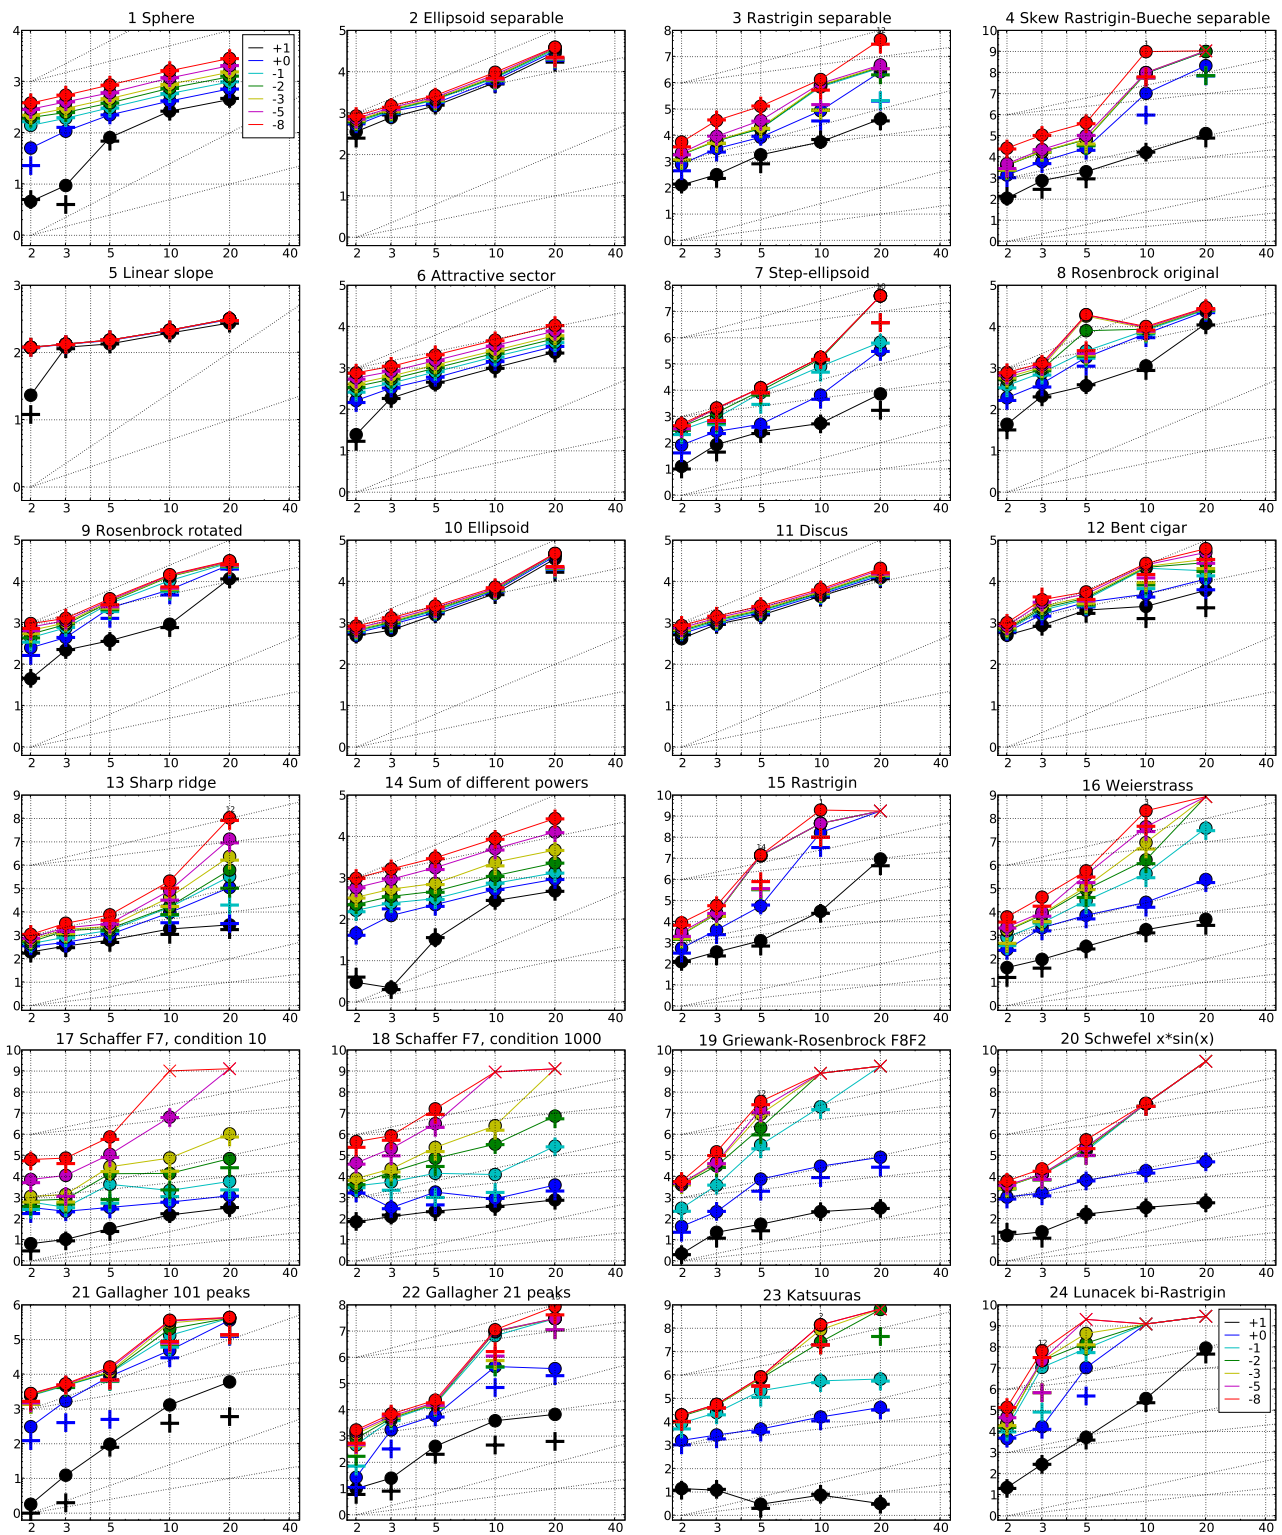


Figure 5: Expected Running Time (ERT, ●) to reach $f_{\text{opt}} + \Delta f$ and median number of function evaluations of successful trials (+), shown for $\Delta f = 10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-5}, 10^{-8}$ (the exponent is given in the legend of f_1 and f_{24}) versus dimension in log-log presentation. The ERT(Δf) equals to $\#FES(\Delta f)$ divided by the number of successful trials, where a trial is successful if $f_{\text{opt}} + \Delta f$ was surpassed during the trial. The $\#FES(\Delta f)$ are the total number of function evaluations while $f_{\text{opt}} + \Delta f$ was not surpassed during the trial from all respective trials (successful and unsuccessful), and f_{opt} denotes the optimal function value. Crosses (×) indicate the total number of function evaluations $\#FES(-\infty)$. Numbers above ERT-symbols indicate the number of successful trials. Annotated numbers on the ordinate are decimal logarithms. Additional grid lines show linear and quadratic scaling.

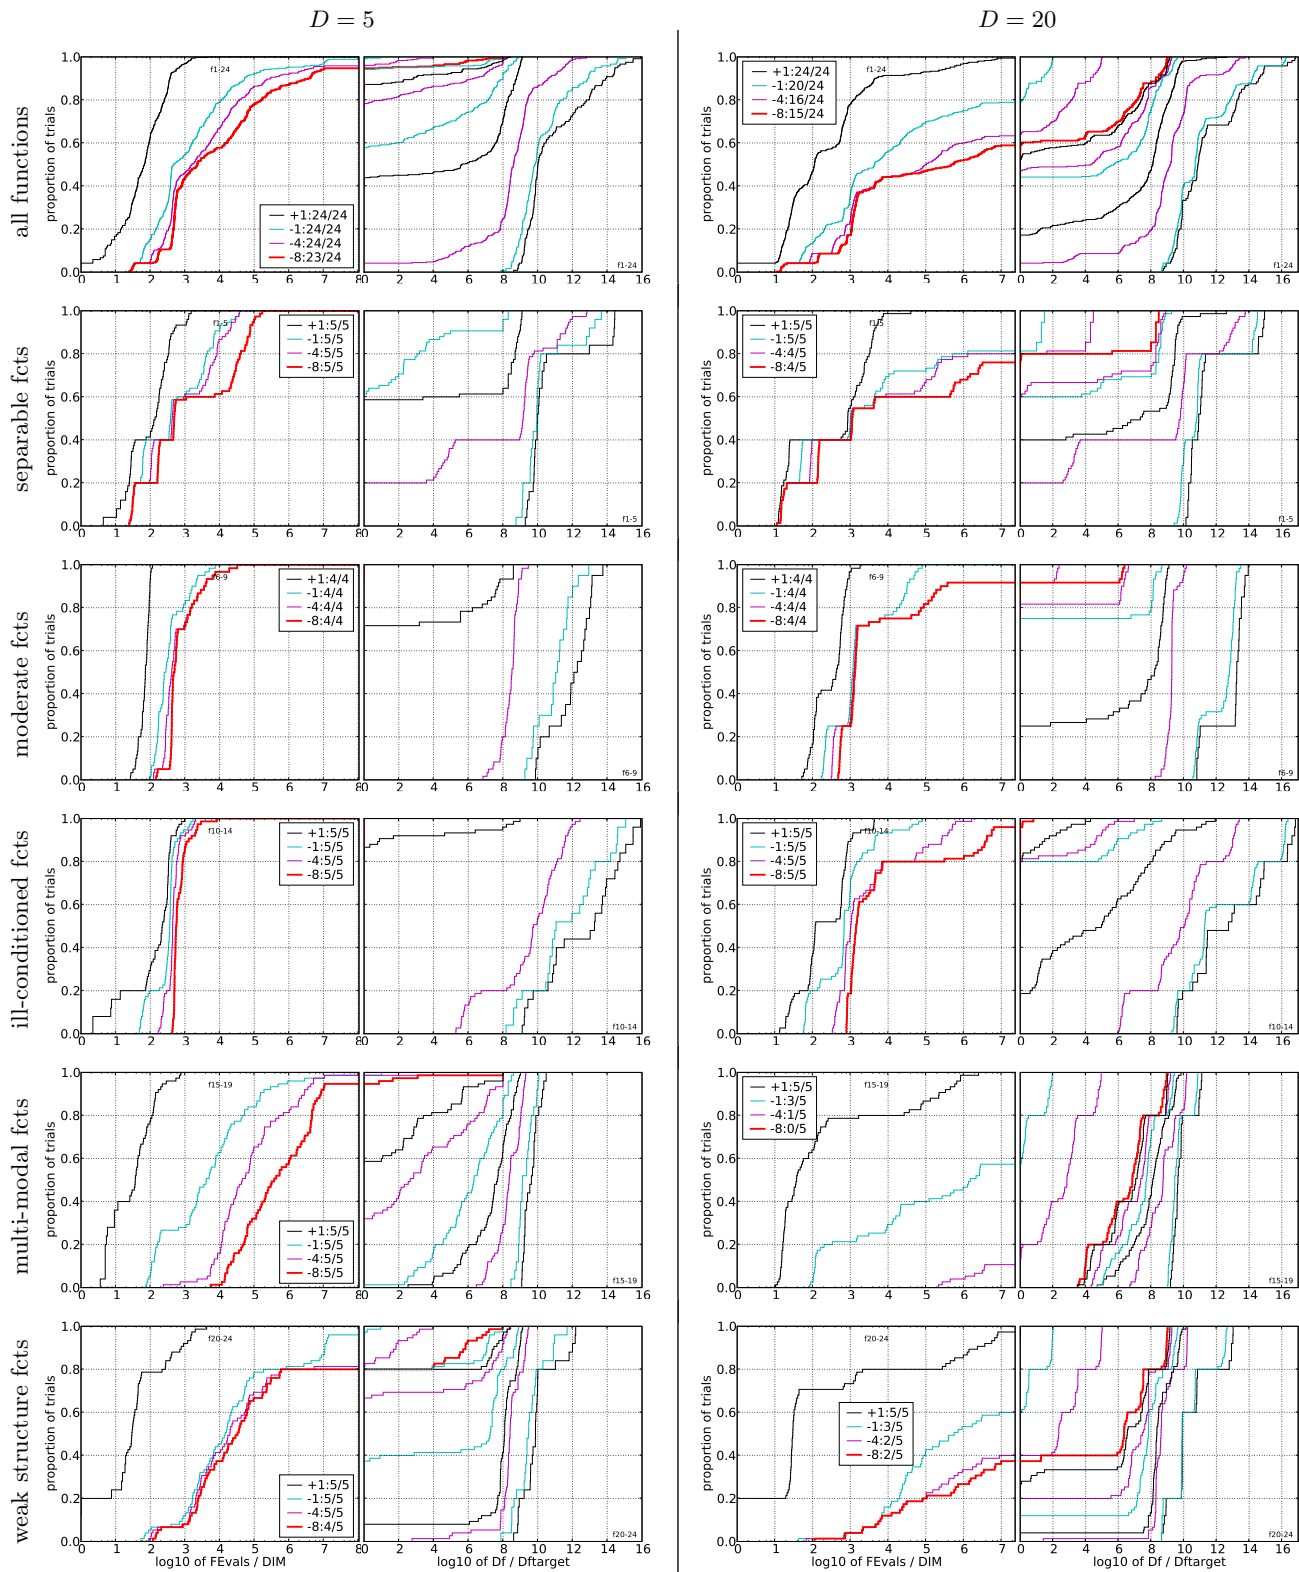


Figure 6: Empirical cumulative distribution functions (ECDFs), plotting the fraction of trials versus running time (left) or Δf . Left subplots: ECDF of the running time (number of function evaluations), divided by search space dimension D , to fall below $f_{\text{opt}} + \Delta f$ with $\Delta f = 10^k$, where k is the first value in the legend. Right subplots: ECDF of the best achieved Δf divided by 10^k (upper left lines in continuation of the left subplot), and best achieved Δf divided by 10^{-8} for running times of $D, 10D, 100D \dots$ function evaluations (from right to left cycling black-cyan-magenta). Top row: all results from all functions; second row: separable functions; third row: misc. moderate functions; fourth row: ill-conditioned functions; fifth row: multi-modal functions with adequate structure; last row: multi-modal functions with weak structure. The legends indicate the number of functions that were solved in at least one trial. FEvals denotes number of function evaluations, D and DIM denote search space dimension, and Δf and Df denote the difference to the optimal function value.

seconds has passed (according to Figure 2 in [13]). The experiment has been conducted with an Intel(R) Core(TM)2 Quad CPU Q9300 2.5GHz under Fedora release 10 (Cambridge) using the C-code provided and optimisation level 2. The time per function evaluation was 2; 1.5; 1.4; 1.2; 1.7; 3 times 10^{-5} seconds in dimensions 2; 3; 5; 10; 20; 40 respectively.

5. CONCLUSIONS

We have presented a continuous VNS model based on three *specialised EAs*: 1) CMA-ES as an EA specialised in generating a good starting point, as generation component, 2) Continuous Local EA, specialised in exploiting local information, as improvement component, and 3) μ CHC, which provides local diversity, as shaking component. Experiments have been carried out on the noiseless Black-Box Optimization Benchmark 2009 testbed.

6. ACKNOWLEDGMENTS

This work was supported by Research Projects TIN2008-05854 and P08-TIC-4173.

7. REFERENCES

- [1] A. Auger and N. Hansen. Performance evaluation of an advanced local search evolutionary algorithm. In D. Corne et al., editors, *Proc. of the IEEE Int. Conf. on Evolutionary Computation*, volume 2, pages 1777–1784. IEEE Press, 2005.
- [2] I. Charon and O. Hudry. The noising methods: A generalization of some metaheuristics. *European Journal of Operational Research*, 135(1):86–101, 2001.
- [3] C. Coello, G. Lamont, and D. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer, 2007.
- [4] K. Deb, A. Anand, and D. Joshi. A computationally efficient evolutionary algorithm for real-parameter optimization. *Evolutionary Computation*, 10(4):371–395, 2002.
- [5] A. Eiben and J. Smith. *Introduction to Evolutionary Computing*. Springer-Verlag, 2003.
- [6] L. Eshelman and J. Schaffer. Real-coded genetic algorithms and interval-schemata. In L. Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 187–202. Morgan Kaufmann, 1993.
- [7] C. Fernandes and A. Rosa. A study on non-random mating and varying population size in genetic algorithms using a royal road function. In *Congress on Evolutionary Computation*, pages 60–66, 2001.
- [8] S. Finck, N. Hansen, R. Ros, and A. Auger. Real-parameter black-box optimization benchmarking 2009: Presentation of the noiseless functions. Technical Report 2009/20, Research Center PPE, 2009.
- [9] C. García-Martínez and M. Lozano. Local search based on genetic algorithms. In P. Siarry and Z. Michalewicz, editors, *Advances in Metaheuristics for Hard Optimization*, Natural Computing, pages 199–221. Springer, 2008.
- [10] C. García-Martínez, M. Lozano, F. Herrera, D. Molina, and A. Sánchez. Global and local real-coded genetic algorithms based on parent-centric crossover operators. *European Journal of Operational Research*, 185(3):1088–1113, 2008.
- [11] C. García-Martínez, M. Lozano, and D. Molina. A local genetic algorithm for binary-coded problems. In T. Runarsson et al., editors, *Proc. of the Int. Conf. on Parallel Problem Solving from Nature*, volume 4193 of *LNCS*, pages 192–201. Springer, 2006.
- [12] F. Glover and G. Kochenberger, editors. *Handbook of Metaheuristics*. Kluwer Academic Publishers, 2003.
- [13] N. Hansen, A. Auger, S. Finck, and R. Ros. Real-parameter black-box optimization benchmarking 2009: Experimental setup. Technical Report RR-6828, INRIA, 2009.
- [14] N. Hansen, S. Finck, R. Ros, and A. Auger. Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions. Technical Report RR-6829, INRIA, 2009.
- [15] N. Hansen and S. Kern. Evaluating the CMA evolution strategy on multimodal test functions. In X. Yao et al., editors, *Int. Conf. on Parallel Problem Solving from Nature*, volume 3242 of *LNCS*, pages 282–291. Springer Berlin, Heidelberg, 2004.
- [16] N. Hansen, S. Müller, and P. Koumoutsakos. Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18, 2003.
- [17] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.
- [18] F. Herrera, M. Lozano, and J. Verdegay. Tackling real-coded genetic algorithms: operators and tools for behavioral analysis. *Artificial Intelligence Reviews*, 12(4):265–319, 1998.
- [19] K. Katayama and H. Narihisa. A new iterated local search algorithm using genetic crossover for the travelling salesman problem. In *ACM Symposium on Applied Computing*, pages 302–306, 1999.
- [20] S. Kirkpatrick, C. Gelatt Jr, and M. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [21] S. Kukkonen and K. Deb. A fast and effective method for pruning of non-dominated solutions in many-objective problems. In *Proc. of the Int. Conf. on Parallel Problem Solving from Nature*, volume 4193 of *LNCS*, pages 553–562, 2006.
- [22] M. Lozano and C. García-Martínez. An evolutionary ILS-perturbation technique. In *Proc. of the Int. Workshop on Hybrid Metaheuristics*, volume 5296 of *LNCS*, pages 1–15, 2008.
- [23] M. Lozano and C. García-Martínez. Hybrid metaheuristics with evolutionary algorithms specializing in intensification and diversification: Overview and progress report. *Computers & Operations Research*. In press, 2009.
- [24] M. Lozano, F. Herrera, N. Krasnogor, and D. Molina. Real-coded memetic algorithms with crossover hill-climbing. *Evolutionary Computation*, 12(3):273–302, 2004.
- [25] N. Mladenovic and P. Hansen. Variable neighborhood search. *Computers & Operations Research*, 24:1097–1100, 1997.
- [26] E. Talbi. A taxonomy of hybrid metaheuristics. *Journal of Heuristics*, 8(5):541–564, 2002.