

Benchmarking the Borg MOEA on the Biobjective bbob-biobj Testbed

Dimo Brockhoff, Pascal Capetillo, Jonathan
Hornwall, Raphael Walker



Introduction

Evolutionary Computing Algorithms

Components of the Borg MOEA

- ϵ -archive
- Randomized restart
- Auto-adaptive multi-operator recombination

Adaptation of the ϵ

Results

Conclusion



Evolutionary Computing Algorithms

- The goal is to find the pareto front

Definition 2.1. For $\epsilon \in \mathbb{R}^+$, a vector $u \in \mathbb{R}^n$ Pareto-dominates another vector $v \in \mathbb{R}^n$ if $\forall i \in \{1, 2, \dots, n\} (u_i \leq v_i)$ and $\exists i \in \{1, 2, \dots, n\} (u_i < v_i)$. This is denoted $v > u$

Definition 2.2. For $\epsilon \in \mathbb{R}^+$, a vector $u \in \mathbb{R}^n$ ϵ -dominates another vector $v \in \mathbb{R}^n$ if $\forall i \in \{1, 2, \dots, n\} (u_i \leq v_i + \epsilon)$ and $\exists i \in \{1, 2, \dots, n\} (u_i < v_i + \epsilon)$

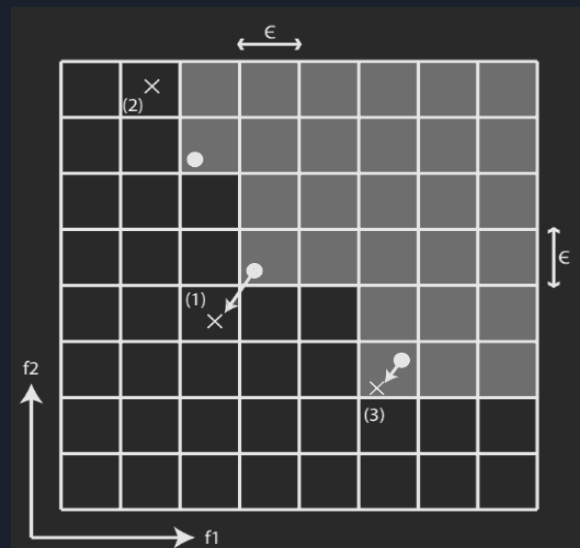
- A population of solutions are “evolved” via evolutionary operators to explore the solution space.
- A record of all non-pareto dominated solutions is kept in an archive.
- Over time, the archive will (hopefully) approximate the pareto front.

Components of the Borg: ϵ -archive

Epsilon archive

Epsilon boxes

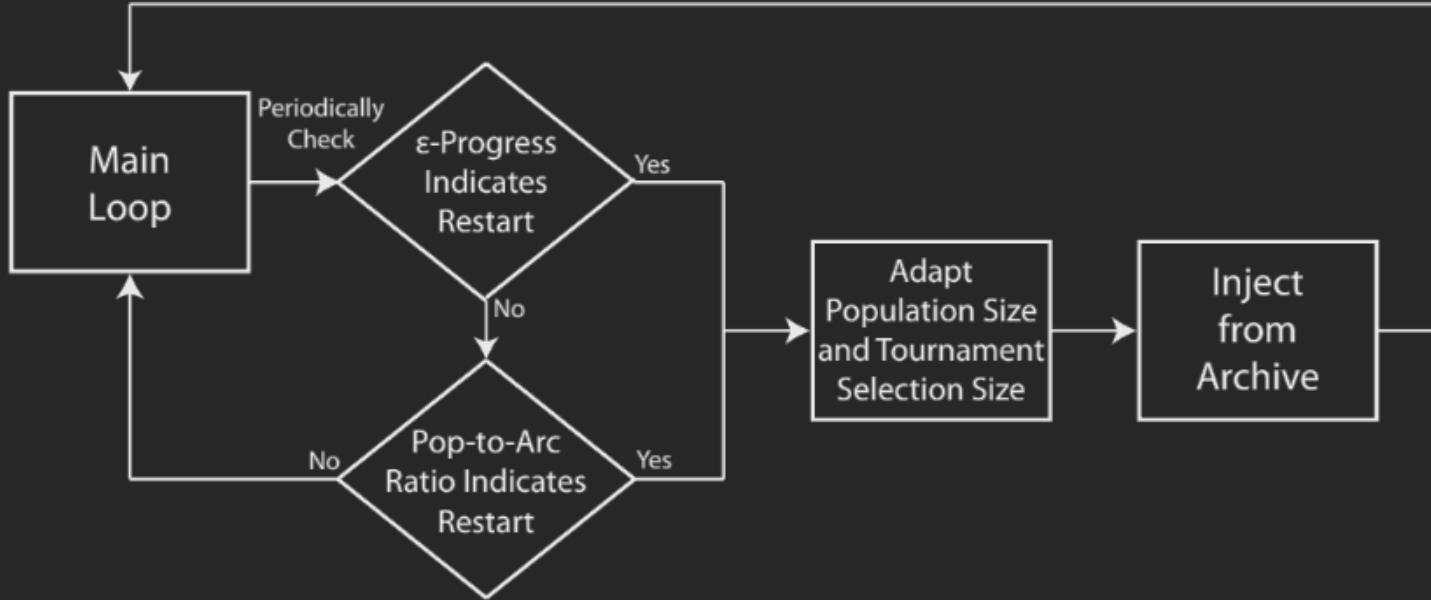
Stagnation



Definition 2.3. For $\epsilon \in \mathbb{R}^+$, a vector $u \in \mathbb{R}^n$ ϵ -box dominates another vector $v \in \mathbb{R}^n$ if one of the following holds


- $\lfloor u/\epsilon \rfloor \leq \lfloor v/\epsilon \rfloor$
- $\lfloor u/\epsilon \rfloor = \lfloor v/\epsilon \rfloor$ and $\|u - \lfloor u/\epsilon \rfloor\| < \|v - \lfloor v/\epsilon \rfloor\|$

Components of the Borg: Randomized restart




Components of the Borg: Auto-adaptive multi-operator recombination


Simulated Binary Crossover




Differential Evolution




Uniform Mutation




Parent-Centric Crossover



Unimodal Normal Distribution Crossover



Simplex Crossover


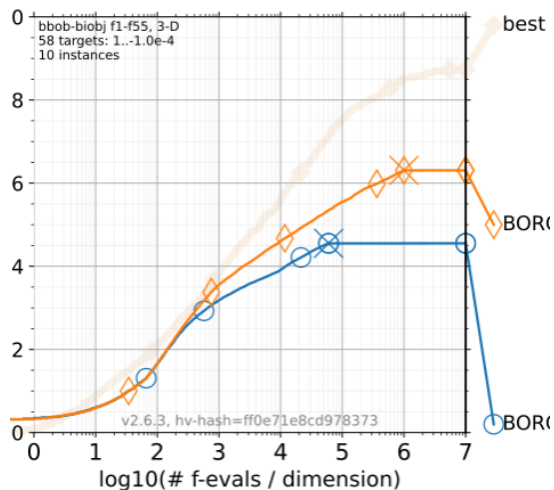







Adaptation of the ϵ

- The ϵ parameter associated to a certain objective is scaled to the range of that objective
- The range of an objective is approximated as the algorithm runs

Results



Results





Conclusion

- Does not outperform current top-of-class algorithms on the bbob biobj suite
- Test suite does not account for more than two objectives, a key use-case for the Borg MOEA
- Parameter tuning of the ϵ has positive impact on the performance. Further tuning an avenue for continued research

Thanks for listening!

