



**HAL**  
open science

# A Mathematical Runtime Analysis of the Non-dominated Sorting Genetic Algorithm III (NSGA-III)

Simon Wietheger, Benjamin Doerr

► **To cite this version:**

Simon Wietheger, Benjamin Doerr. A Mathematical Runtime Analysis of the Non-dominated Sorting Genetic Algorithm III (NSGA-III). GECCO 2024 Proceedings Companion, 2024, Melbourne, Australia. pp.63 - 64, 10.1145/3638530.3664062 . hal-04676955

**HAL Id: hal-04676955**

**<https://hal.science/hal-04676955>**

Submitted on 24 Aug 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



# Hot off the Press: A Mathematical Runtime Analysis of the Non-dominated Sorting Genetic Algorithm III (NSGA-III)

Simon Wietheger

Algorithms and Complexity Group  
TU Wien  
Vienna, Austria

Benjamin Doerr

Laboratoire d'Informatique (LIX)  
École Polytechnique, CNRS  
Institut Polytechnique de Paris  
Palaiseau, France

## ABSTRACT

The Non-dominated Sorting Genetic Algorithm II (NSGA-II) is the most prominent multi-objective evolutionary algorithm for real-world applications. While it performs evidently well on bi-objective optimization problems, empirical studies suggest that it is less effective when applied to problems with more than two objectives. A recent mathematical runtime analysis confirmed this observation by proving that the NSGA-II for an exponential number of iterations misses a constant factor of the Pareto front of the simple  $m$ -objective ONEMINMAX problem when  $m \geq 3$ .

In this work, we provide the first mathematical runtime analysis of the NSGA-III, a refinement of the NSGA-II aimed at better handling more than two objectives. We prove that the NSGA-III with sufficiently many reference points – a small constant factor more than the size of the Pareto front, as suggested for this algorithm – computes the complete Pareto front of the 3-objective ONEMINMAX benchmark in an expected number of  $O(n \log n)$  iterations. This result holds for all population sizes (that are at least the size of the Pareto front). It shows a drastic advantage of the NSGA-III over the NSGA-II on this benchmark.

This paper for the Hot-off-the-Press track at GECCO 2024 summarizes the work *Simon Wietheger and Benjamin Doerr. A mathematical runtime analysis of the Non-dominated Sorting Genetic Algorithm III (NSGA-III). In International Joint Conference on Artificial Intelligence, IJCAI 2023. 5657–5665, 2023. [15]*.

## CCS CONCEPTS

• **Theory of computation** → *Theory of randomized search heuristics.*

## KEYWORDS

NSGA-III, runtime analysis, multi-objective optimization, theory.

## ACM Reference Format:

Simon Wietheger and Benjamin Doerr. 2024. Hot off the Press: A Mathematical Runtime Analysis of the Non-dominated Sorting Genetic Algorithm III (NSGA-III). In *Genetic and Evolutionary Computation Conference (GECCO '24 Companion)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3638530.3664062>

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).  
*GECCO '24 Companion, July 14–18, 2024, Melbourne, VIC, Australia*  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0495-6/24/07.  
<https://doi.org/10.1145/3638530.3664062>

## SUMMARY OF OUR RESULTS

Many practical applications require to optimize for multiple, conflicting objectives. Such tasks can be tackled by population-based algorithms, whose population eventually represents a set of Pareto solutions, solutions that cannot strictly be dominated by any other solution. Thereby, they represent multiple useful trade-offs between the objectives and allow the user to choose among these according to their personal preferences. Indeed, evolutionary algorithms (EAs), or, more precisely, multi-objective evolutionary algorithms (MOEAs), have been successfully applied to many real-world problems [20]. Among these, Zhou et al. [20] identify the non-dominated sorting genetic algorithm (NSGA-II) [6] as the most prominent one. Both empirical evaluations [11, 13] and recent mathematical runtime analyses [1–4, 7–9, 19] confirm the strong results of the NSGA-II on bi-objective benchmarks. The performance on problems with 3 or more objectives, however, is not as well understood. Empirical studies, for example [11], suggest that the NSGA-II struggles with such problems. A recent mathematical runtime analysis [17] shows that the NSGA-II regularly loses desirable solutions when optimizing the  $m$ -objective ONEMINMAX problem for  $m \geq 3$ , and consequently, cannot find its Pareto front (the set of Pareto optimal solution values) in sub-exponential time. As a remedy, Deb and Jain [5] proposed a modified version of the NSGA-II, called NSGA-III. It replaces the crowding distance, a measure which the NSGA-II uses in addition to the dominance relation to determine which individuals are taken in the next generation, by a procedure involving reference points in the solution space. Their evaluations on benchmarks with 3 to 15 objectives show that the NSGA-III is suitable for more than 2 objectives.

These empirical insights are, however, not yet supported by a deeper mathematical understanding of the working principles of the NSGA-III. In order to fill this gap, we mathematically analyze the runtime of the NSGA-III on the 3-objective ONEMINMAX (3OMM) problem. The first objective of this 3-objective benchmark is the total number of 0-bits, while the second and third objective are the numbers of 1-bits in the first and second half of the bitstring, respectively. We show that by employing sufficiently many reference points (a small constant factor more than the size of the Pareto front, as suggested for this algorithm) and a population at least of the size of the Pareto front,  $N \geq (\frac{n}{2} + 1)^2$ , once a solution for a point on the Pareto front is found, the population will always contain such a solution. This is a notable difference to the NSGA-II [17] and enables us to prove that after an expected number of  $O(n \log n)$  iterations the NSGA-III (for all future iterations) has a population that covers the Pareto front. Overall, this result indicates, in a rigorous manner, that the selection mechanism of the

NSGA-III has significant advantages over the one of the NSGA-II. Possibly, our result also indicates that more algorithm users should switch from the NSGA-II, still the dominant algorithm in practice, to the NSGA-III. We note that the latter has as additional parameter the number of reference points, but the general recommendation to use by a small factor more reference points than the size of the Pareto front (or, in the case of approximate solutions, the size of the desired solution set) renders it easy to choose this parameter. We also note that our results support this parameter choice, our proven guarantees also hold from the point on when the number of reference points is a small constant factor larger than the Pareto front. We finally note that using more reference points does not significantly increase the runtime (not at all when counting fitness evaluations and only moderately when counting wall-clock time), so in any case the choice of this parameter appears not too critical.

In a small set of experiments, we demonstrate that the NSGA-III is indeed able to solve 3OMM while the NSGA-II, even with larger population sizes, struggles to cover more than a certain ratio of the Pareto front. Further, our experiments show that the runtime of the NSGA-III is increasing when using crossover in addition to the mutation operator. This suggests that progress is made mainly due to mutation, and justifies that our analyses do not examine the potential benefits of crossover.

*Summary and conclusion:* We provide a theoretical understanding on how the NSGA-III copes with three objectives where the NSGA-II struggles: using reference points instead of crowding distance as a tie breaker in the survival process guarantees that non-dominated solutions are never lost in future iterations (for suitable choices of the parameters). Using this insight, we give a first mathematical runtime analysis of the NSGA-III and show that it solves the 3OMM benchmark in  $O(n \log(n))$  iterations (when using a sufficiently large population size and sufficiently many reference points).

*Subsequent work:* Our analysis of optimization time of the NSGA-III has been extended in several directions. Most importantly, the work of Opris, Dang, Neumann and Sudholt [12] generalizes our central result of not losing values on the Pareto front to general problems with an arbitrary, but constant, number of objectives. They use this insight to analyze the optimization time of the NSGA-III on three established many-objective benchmarks. In our independent and parallel study of the optimization time of various multi-objective evolutionary algorithms and on many-objective benchmarks (also for non-constant number of objectives) [16], we argued that these results extend to the NSGA-III provided that a result as shown in [12] is shown (hence now our results also hold for the NSGA-III). Based on our paper, Qian [14] suggests to study the NSGA-III for approximating clusterings. Zheng and Doerr [18] state that our results just like their analysis of the SMS-EMOA (which is very similar to the NSGA-II but also replaces the crowding distance tie breaker) show that the struggles of the NSGA-II with more than 2 objectives are not due to its general structure but are caused by using crowding distance as a tie breaker. A similar conclusion can be drawn from [10], who show that the NSGA-II with a further tie-breaker after the crowding distance can easily optimize several classic many-objective benchmarks.

## ACKNOWLEDGMENTS

Simon Wietheger acknowledges the support by a fellowship via the International Exchange Program of École Polytechnique and by the Austrian Science Fund (FWF, START Project Y1329). This research benefited from the support of the FMJH Program Gaspard Monge for optimization and operations research and their interactions with data science.

## REFERENCES

- [1] Chao Bian and Chao Qian. 2022. Better running time of the non-dominated sorting genetic algorithm II (NSGA-II) by using stochastic tournament selection. In *Parallel Problem Solving From Nature, PPSN 2022*. Springer, 428–441.
- [2] Sacha Cerf, Benjamin Doerr, Benjamin Hebras, Jakob Kahane, and Simon Wietheger. 2023. The first proven performance guarantees for the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) on a combinatorial optimization problem. In *International Joint Conference on Artificial Intelligence, IJCAI 2023*. ijcai.org, 5522–5530.
- [3] Duc-Cuong Dang, Andre Opris, Bahare Salehi, and Dirk Sudholt. 2023. Analysing the robustness of NSGA-II under noise. In *Genetic and Evolutionary Computation Conference, GECCO 2023*. ACM, 642–651.
- [4] Duc-Cuong Dang, Andre Opris, Bahare Salehi, and Dirk Sudholt. 2023. A proof that using crossover can guarantee exponential speed-ups in evolutionary multi-objective optimisation. In *Conference on Artificial Intelligence, AAAI 2023*. AAAI Press, 12390–12398.
- [5] Kalyanmoy Deb and Himanshu Jain. 2014. An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints. *IEEE Transactions on Evolutionary Computation* 18 (2014), 577–601.
- [6] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6 (2002), 182–197.
- [7] Benjamin Doerr and Zhongdi Qu. 2023. A first runtime analysis of the NSGA-II on a multimodal problem. *IEEE Transactions on Evolutionary Computation* 27 (2023), 1288–1297.
- [8] Benjamin Doerr and Zhongdi Qu. 2023. From understanding the population dynamics of the NSGA-II to the first proven lower bounds. In *Conference on Artificial Intelligence, AAAI 2023*. AAAI Press, 12408–12416.
- [9] Benjamin Doerr and Zhongdi Qu. 2023. Runtime analysis for the NSGA-II: Provable speed-ups from crossover. In *Conference on Artificial Intelligence, AAAI 2023*. AAAI Press, 12399–12407.
- [10] Tudor Ivan, Martin S. Krejca, and Benjamin Doerr. 2024. Speeding Up the NSGA-II With a Simple Tie-Breaking Rule. (2024). Preprint.
- [11] Vineet Khare, Xin Yao, and Kalyanmoy Deb. 2003. Performance Scaling of Multi-objective Evolutionary Algorithms. In *Evolutionary Multi-Criterion Optimization (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 376–390.
- [12] Andre Opris, Duc Cuong Dang, Frank Neumann, and Dirk Sudholt. 2024. Runtime analyses of NSGA-III on many-objective problems. In *Genetic and Evolutionary Computation Conference, GECCO 2024*. ACM. To appear.
- [13] Robin C. Purshouse and Peter J. Fleming. 2007. On the Evolutionary Optimization of Many Conflicting Objectives. *IEEE Transactions on Evolutionary Computation* 11 (2007), 770–784. <https://doi.org/10.1109/TEVC.2007.910138>
- [14] Chao Qian. 2023. Can Evolutionary Clustering Have Theoretical Guarantees? *IEEE Transactions on Evolutionary Computation* (2023), 1–1. <https://doi.org/10.1109/TEVC.2023.3296645>
- [15] Simon Wietheger and Benjamin Doerr. 2023. A mathematical runtime analysis of the Non-dominated Sorting Genetic Algorithm III (NSGA-III). In *International Joint Conference on Artificial Intelligence, IJCAI 2023*. ijcai.org, 5657–5665.
- [16] Simon Wietheger and Benjamin Doerr. 2024. Near-Tight Runtime Guarantees for Many-Objective Evolutionary Algorithms. (2024). Preprint.
- [17] Weijie Zheng and Benjamin Doerr. 2023. Runtime analysis for the NSGA-II: proving, quantifying, and explaining the inefficiency for many objectives. *IEEE Transactions on Evolutionary Computation* (2023). In press, <https://doi.org/10.1109/TEVC.2023.3320278>.
- [18] Weijie Zheng and Benjamin Doerr. 2024. Runtime analysis of the SMS-EMOA for many-objective optimization. In *Conference on Artificial Intelligence, AAAI 2024*. AAAI Press.
- [19] Weijie Zheng, Yufei Liu, and Benjamin Doerr. 2022. A first mathematical runtime analysis of the Non-Dominated Sorting Genetic Algorithm II (NSGA-II). In *Conference on Artificial Intelligence, AAAI 2022*. AAAI Press, 10408–10416.
- [20] Aimin Zhou, Bo-Yang Qu, Hui Li, Shi-Zheng Zhao, Ponnuthurai Nagarathnam Suganthan, and Qingfu Zhang. 2011. Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation* 1 (2011), 32–49.