



HAL
open science

Hot off the Press: Runtime Analysis for the NSGA-II: Proving, Quantifying, and Explaining the Inefficiency For Many Objectives

Weijie Zheng, Benjamin Doerr

► **To cite this version:**

Weijie Zheng, Benjamin Doerr. Hot off the Press: Runtime Analysis for the NSGA-II: Proving, Quantifying, and Explaining the Inefficiency For Many Objectives. GECCO '24 Companion: Genetic and Evolutionary Computation Conference Companion, 2024, Melbourne, Australia. pp.67-68, 10.1145/3638530.3664061 . hal-04676966

HAL Id: hal-04676966

<https://hal.science/hal-04676966>

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: Runtime Analysis for the NSGA-II: Proving, Quantifying, and Explaining the Inefficiency For Many Objectives

Weijie Zheng

School of Computer Science and Technology
International Research Institute for Artificial Intelligence
Harbin Institute of Technology
Shenzhen, China

Benjamin Doerr

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

ABSTRACT

The NSGA-II is one of the most prominent algorithms to solve multi-objective optimization problems. Despite numerous successful applications, several studies have shown that the NSGA-II is less effective for larger numbers of objectives. In this work, we use mathematical runtime analyses to rigorously demonstrate and quantify this phenomenon. We show that even on the simple m -objective generalization of the discrete OneMinMax benchmark, where every solution is Pareto optimal, the NSGA-II also with large population sizes cannot compute the full Pareto front (objective vectors of all Pareto optima) in sub-exponential time when the number of objectives is at least three. The reason for this unexpected behavior lies in the fact that in the computation of the crowding distance, the different objectives are regarded independently. This is not a problem for two objectives, where any sorting of a pairwise incomparable set of solutions according to one objective is also such a sorting according to the other objective (in the inverse order).

This paper for the Hot-off-the-Press track at GECCO 2024 summarizes the work *Weijie Zheng, Benjamin Doerr: Runtime Analysis for the NSGA-II: Proving, Quantifying, and Explaining the Inefficiency For Many Objectives*. *IEEE Transactions on Evolutionary Computation*, in press. <https://doi.org/10.1109/TEVC.2023.3320278> [23].

CCS CONCEPTS

• **Theory of computation** → **Theory and algorithms for application domains**; *Theory of randomized search heuristics*.

KEYWORDS

NSGA-II, runtime analysis, many-objective optimization, theory

ACM Reference Format:

Weijie Zheng and Benjamin Doerr. 2024. Hot off the Press: Runtime Analysis for the NSGA-II: Proving, Quantifying, and Explaining the Inefficiency For Many Objectives. 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/10.1145/3638530.3664061>

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/10.1145/3638530.3664061>

SUMMARY OF OUR RESULTS

The most prominent multi-objective evolutionary algorithm (MOEA) with many successful applications in various domains is the *non-dominated sorting genetic algorithm II* (NSGA-II) [5]. Despite numerous positive results on the NSGA-II, several studies have observed that the NSGA-II is less effective when the number of objectives grows. From their experimental comparison of three MOEAs, Khare, Yao, and Deb [11] reported that the NSGA-II increasingly suffers in converging to the global Pareto front when the number of objectives increases from 2 to 8. This early study is a pure performance comparison, so no explanations for the different behaviors were found. Also, it should be noted that the other algorithms regarded also suffered in different ways from growing numbers of objectives. Purshouse and Fleming [15] aimed for further empirical observations. They studied the range of suitable configurations of the NSGA-II that can result in good solutions. That the suitable ranges shrink along with the increasing number of objectives, indicates the increasing difficulty for the NSGA-II for more objectives. They also observed some factors that might be relevant for the poor performance of the NSGA-II, like that when the population evolves, the proportion of non-dominated solutions in the population rapidly increases to 100%, and that in each iteration, only a low proportion of the newly generated solutions dominates the current solutions. These observations are regarded as one of the key challenges for many-objective optimizations, see the surveys [9, 13].

Since apparently there is a lack of understanding of the performance of the NSGA-II for many-objective problems, we try to approach this research question via a mathematical runtime analysis. Such analyses are an integral part of the theory of heuristic search. While often restricted to simple algorithmic settings, this alternative approach has led to several deep and very reliable (namely mathematically proven) results in the past. Also, often the proofs also reveal the reason why a certain phenomenon can be observed.

Our Contribution: To this aim, we conduct a mathematical runtime analysis of the NSGA-II on the m -objective version of the classic bi-objective ONEMINMAX benchmark. This pseudo-Boolean (that is, defined on bit-strings of length n) benchmark is very simple in several respects, for example, any solution is Pareto optimal and the objectives are all equivalent to the ONEMAX benchmark, which is generally considered as the easiest single-objective Pseudo-boolean benchmark [6, 16, 20]. For the bi-objective ONEMINMAX problem, a good performance of the NSGA-II has been proven recently [2, 21, 22, 25].

Our main (proven) result is that – despite the simplicity of the problem and in drastic contrast to the bi-objective setting – for all

numbers $m \geq 3$ of objectives the NSGA-II also with large population sizes cannot find the full Pareto front faster than in exponential time. Even worse, we prove that for an exponential time, the population of the NSGA-II will miss a constant fraction of the Pareto front. Our experiments confirm this finding in a very clear manner.

Our mathematical runtime analysis also gives an explanation for the drastic change of behavior between two and three objectives, namely that in the definition of the crowding distance the different objectives are regarded independently. This is not a problem for two objectives, because for a set of pair-wise non-dominated solutions a sorting with respect to one objective automatically is a sorting with respect to the other objective (in the opposite order). Hence here in fact the two objectives automatically are not treated independently. From three objectives on, such a correlation between the objectives does not exist, and this can lead to the problems made precise in this work.

This understanding has two implications. On the negative side, it appears very likely that the difficulties observed for the ONEMINMAX benchmark will also occur for many other optimization problems, including continuous optimization problems. On the positive side, this understanding suggests to search for an alternative crowding distance measure that does not treat the objectives independently.

We note that given the known difficulties of the NSGA-II with more objectives, several alternatives to the crowding distance have been proposed. For example, Deb and Jain [3] proposed the NSGA-III, where the crowding distance is replaced by a system of reference points. In their SMS-EMOA, Beume, Naujoks, and Emmerich [1] replaced the crowding distance by the hypervolume. The hypervolume measure was also considered as the second sorting criterion in the multiobjective CMA-ES [8]. It was also used in hybrid with the crowding distance [18]. Instead of removing points with the smallest crowding distances in the NSGA-II, the diversity preservation way via clustering from SPEA [26] was used to select the survived individuals [4, 17]. However, none of these alternatives is as accepted in practice as the NSGA-II. For this reason, we hope that our work can not only spur the development of superior algorithms, but also motivate practitioners to try moving from the NSGA-II to these more modern algorithms. We further note that several promising variations of the crowding distance have been proposed, e.g., [7, 12], which again are not used a lot in practice.

Subsequent works: Two recent runtime analyses [14, 19] revealed that the difficulties proven here for the NSGA-II do not arise with the NSGA-III. In [24] it was proven that the SMS-EMOA can overcome these difficulties as well. Very recently [10], it was shown that the NSGA-II with an additional tie-breaker after the crowding distance, similar to the one proposed in [7], can optimize the many-objective versions of several classic benchmarks efficiently when the population size is slightly larger than the size of the Pareto front. However, since this variant only differs from the traditional NSGA-II when objective values occur multiple times in the population, we see that this variant displays the same unfavorable behavior on Example 2 in [23] as the traditional NSGA-II.

REFERENCES

[1] Nicola Beume, Boris Naujoks, and Michael Emmerich. 2007. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of*

Operational Research 181 (2007), 1653–1669.

[2] 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.

[3] 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.

[4] Kalyanmoy Deb, Manikanth Mohan, and Sikhar Mishra. 2003. A fast multi-objective evolutionary algorithm for finding well-spread Pareto-optimal solutions. *KanGAL report 2003002* (2003).

[5] 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.

[6] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. 2012. Multiplicative drift analysis. *Algorithmica* 64 (2012), 673–697.

[7] Félix-Antoine Fortin and Marc Parizeau. 2013. Revisiting the NSGA-II crowding-distance computation. In *Genetic and Evolutionary Computation Conference, GECCO 2013*. ACM, 623–630.

[8] Christian Igel, Nikolaus Hansen, and Stefan Roth. 2007. Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation* 15 (2007), 1–28.

[9] Hisao Ishibuchi, Noritaka Tsukamoto, and Yusuke Nojima. 2008. Evolutionary many-objective optimization: A short review. In *IEEE Congress on Evolutionary Computation, CEC 2008*. IEEE, 2419–2426.

[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 *International Conference on Evolutionary Multi-criterion Optimization, EMO 2003*. Springer, 376–390.

[12] Saku Kukkonen and Kalyanmoy Deb. 2006. Improved pruning of non-dominated solutions based on crowding distance for bi-objective optimization problems. In *Conference on Evolutionary Computation, CEC 2006*. IEEE, 1179–1186.

[13] Bingdong Li, Jinlong Li, Ke Tang, and Xin Yao. 2015. Many-objective evolutionary algorithms: A survey. *ACM Computing Surveys (CSUR)* 48 (2015), 1–35.

[14] 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.

[15] 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.

[16] Dirk Sudholt. 2013. A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation* 17 (2013), 418–435.

[17] Vimal L. Vachhani, Vipul K. Dabhi, and Harshadkumar B. Prajapati. 2016. Improving NSGA-II for solving multi objective function optimization problems. In *International Conference on Computer Communication and Informatics, ICCCI 2016*. IEEE, 1–6.

[18] Yali Wang, Michael Emmerich, André Deutz, and Thomas Bäck. 2019. Diversity-indicator based multi-objective evolutionary algorithm: DI-MOEA. In *International Conference on Evolutionary Multi-Criterion Optimization, EMO 2019*. Springer, 346–358.

[19] 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.

[20] Carsten Witt. 2013. Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Combinatorics, Probability & Computing* 22 (2013), 294–318.

[21] Weijie Zheng and Benjamin Doerr. 2022. Better approximation guarantees for the NSGA-II by using the current crowding distance. In *Genetic and Evolutionary Computation Conference, GECCO 2022*. ACM, 611–619.

[22] Weijie Zheng and Benjamin Doerr. 2023. Mathematical runtime analysis for the non-dominated sorting genetic algorithm II (NSGA-II). *Artificial Intelligence* 325 (2023), 104016.

[23] 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>.

[24] 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.

[25] 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.

[26] Eckart Zitzler and Lothar Thiele. 1999. Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation* 3 (1999), 257–271.