Hot off the Press: Runtime Analysis of the SMS-EMOA for Many-Objective Optimization

Weijie Zheng, Benjamin Doerr

Hot off the Press: Runtime Analysis of the SMS-EMOA for Many-Objective Optimization

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 was recently proven to have difficulties in many-objective optimization. In contrast, the literature experimentally shows a good performance of the SMS-EMOA, which can be seen as a steady-state NSGA-II that uses the hypervolume contribution instead of the crowding distance as the second selection criterion.

This paper conducts the first rigorous runtime analysis of the SMS-EMOA for many-objective optimization. To this aim, we first propose a many-objective counterpart, the m-objective mOJZJ problem, of the bi-objective OneJumpZeroJump benchmark, which is the first many-objective multimodal benchmark used in a mathematical runtime analysis. We prove that SMS-EMOA computes the full Pareto front of this benchmark in an expected number of $O(M^2 n^k)$ iterations, where n denotes the problem size (length of the bit-string representation), k the gap size (a difficulty parameter of the problem), and $M = (2n/m - 2k + 3)^{m/2}$ the size of the Pareto front. This result together with the existing negative result on the original NSGA-II shows that in principle, the general approach of the NSGA-II is suitable for many-objective optimization, but the crowding distance as tie-breaker has deficiencies.

We obtain three additional insights on the SMS-EMOA. Different from a recent result for the bi-objective OneJumpZeroJump benchmark, the stochastic population update often does not help for mOJZJ. It results in a $1/\Theta(\min\{Mk^{1/2}/2^{k/2}, 1\})$ speed-up, which is $\Theta(1)$ for large m such as $m > k$. On the positive side, we prove that heavy-tailed mutation still results in a speed-up of order $k^{k+0.5-\beta}$. Finally, we conduct the first runtime analyses of the SMS-EMOA on the bi-objective OneMinMax and LOTZ benchmarks and show that it has a performance comparable to the GSEMO and the NSGA-II.

This paper for the Hot-off-the-Press track at GECCO 2024 summarizes the work *Weijie Zheng, Benjamin Doerr: Runtime Analysis of the SMS-EMOA for Many-Objective Optimization. AAAI 2024, 20874–20882* [20].

CCS CONCEPTS

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

GECCO '24 Companion, July 14–18, 2024, Melbourne, VIC, Australia

ACM ISBN 979-8-4007-0495-6/24/07.
https://doi.org/10.1145/3638530.3664064

KEYWORDS

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

ACM Reference Format:
Weijie Zheng and Benjamin Doerr. 2024. Hot off the Press: Runtime Analysis of the SMS-EMOA for Many-Objective Optimization. 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/10.1145/3638530.3664064

SUMMARY OF OUR RESULTS

The NSGA-II is the most widely-applied multiobjective evolutionary algorithm (MOEA). Non-dominated sorting and crowding distance are its two major features differentiating it from basic MOEAs such as the GSEMO or the $(\mu + 1)$ SIBEA. Zheng et al. [22] conducted the first runtime analysis of the NSGA-II (see [17] for the journal version). This work quickly inspired many interesting follow-up results in bi-objective optimization [2, 4–6, 8–10, 16]. In contrast to these positive results for two objectives, Zheng and Doerr [18] proved that for $m \geq 3$ objectives the NSGA-II needs at least exponential time (in expectation and with high probability) to cover the full Pareto front of the m-objective OneMinMax benchmark, a simple many-objective version of the basic OneMax problem where all search points are Pareto optimal. They claimed that the main reason for this low efficiency is the independent computation of the crowding distance in each objective.

A very recent work showed that the NSGA-III, a successor algorithm of the NSGA-II aimed at better coping with many objectives, can efficiently solve the 3-objective OneMinMax problem [14]. Since apparently practitioners much prefer the NSGA-II (more than 17,000 citations on Google scholar only in the last five years) over the NSGA-III (5,035 citations since its publication in 2013), it remains an interesting question whether there are variants of the NSGA-II which better cope with many objectives.

With the SMS-EMOA, an interesting variant of the NSGA-II was proposed by Beume et al. [1]. This algorithm is a steady-state variant of the NSGA-II (that is, in each iteration only a single offspring is generated and possibly integrated into the population) that further replaces the crowding distance as secondary selection criterion with the classic hypervolume contribution. Many empirical works (see the almost 2,000 papers citing [1]) confirmed the good performance of the SMS-EMOA for many-objective optimization. The first mathematical runtime analysis of the SMS-EMOA was conducted very recently by Bian et al. [3], who proved that its expected runtime on the bi-objective OneJumpZeroJump problem is $O(n^{k+1})$. They also proposed a stochastic population update mechanism and proved

that it has the often superior runtime of $O(n^{k+1} \min\{1, n/2^{k/4}\})$. Zheng et al. [21] proved that the SMS-EMOA has an expected runtime of $O(n^4)$ on the bi-objective DLTB problem.

Our Contributions: This paper conducts the first mathematical runtime analysis of the SMS-EMOA for more than two objectives. We first define the mOJZJ benchmark, an m-objective counterpart of the bi-objective OneJumpZeroJump problem [19], which is the problem analyzed in the first runtime analysis for the SMS-EMOA [3]. We note that the mOJZJ problem is the first multimodal many-objective benchmark proposed for a theoretical analysis, to the best of our knowledge. We prove that the SMS-EMOA covers the full Pareto front of this benchmark in an expected number of $O(M^2 n^k)$ iterations, where n is the problems size, k the gap size (a difficulty parameter of the problem), m the number of objectives, and $M = (2n/m - 2k + 3)^{m/2}$ the size of the Pareto front.

We recall that the original NSGA-II needs at least exponential time to optimize the mOneMinMax problem, which is a special case of mOJZJ with gap size $k = 1$. Since the SMS-EMOA employs non-dominated sorting, but replaces the crowding distance in the original NSGA-II by the hypervolume contribution, our result in a similar fashion as the analysis of the NSGA-III in [14] suggests that the general approach of the NSGA-II is suitable for many-objective optimization and that it is only the crowding distance as tie-breaker which is not appropriate for more than two objectives.

We then analyze whether the better performance of the SMS-EMOA on the bi-objective OneJumpZeroJump problem achieved via a new stochastic population update [3] extends to the m-objective mOJZJ problem. Unfortunately, we shall observe that only a speed-up of order $1/\Theta(\min\{Mk^{1/2}/2^{k/2}, 1\})$ is obtained, which is $\Theta(1)$ when m is large, e.g., $m > k$.

On the positive side, we show that the advantage of heavy-tailed mutation is preserved. We analyze the SMS-EMOA with heavy-tailed mutation on the mOJZJ benchmark and prove that a speed-up of order $k^{0.5+k-\beta}$ is achieved. This is the same speed-up as observed for single-objective and bi-objective Jump problems [7, 19]. We note that this is the first theoretical work to support the usefulness of heavy-tailed mutation in many-objective optimization.

Finally, since so far the performance of the SMS-EMOA was only analyzed on the bi-objective OneJumpZeroJump problem [3], we conduct mathematical runtime analyses of the SMS-EMOA also on the two most prominent bi-objective benchmarks OneMinMax and LOTZ. We prove that the SMS-EMOA finds the Pareto fronts of these benchmarks in an expected number of iterations of at most $2e(n+1)n(\ln n + 1)$ for OneMinMax and at most $2en^2(n+1)$ for LOTZ. These are the same asymptotic runtimes (in terms of fitness evaluations) as known for the GSEMO and the NSGA-II.

Subsequent Works: The efficiency of NSGA-III on other classic many-objective problems is shown recently [13]. Both [13] and [15] show improved runtime results for many-objective benchmarks. With our methods, these results extend to the SMS-EMOA. Back to our work's motivation of whether there are variants of NSGA-II which better cope with many objectives, a recent preprint [12] shows that an additional tie-breaker after the crowding distance, similar to the one proposed in [11], will help. However, the difficulty of the original NSGA-II shown in Example 2 in [18] still exists.

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] Chao Bian, Yawen Zhou, Miqing Li, and Chao Qian. 2023. Stochastic population update can provably be helpful in multi-objective evolutionary algorithms. In *International Joint Conference on Artificial Intelligence, IJCAI 2023*. ijcai.org, 5513–5521.

[4] 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.

[5] 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.

[6] 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.

[7] Benjamin Doerr, Huu Phuoc Le, Régis Makhmara, and Ta Duy Nguyen. 2017. Fast genetic algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2017*. ACM, 777–784.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] Tudor Ivan, Martin S. Krejca, and Benjamin Doerr. 2024. Speeding Up the NSGA-II With a Simple Tie-Breaking Rule. (2024). Preprint.

[13] 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.

[14] 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.

[15] Simon Wietheger and Benjamin Doerr. 2024. Near-tight runtime guarantees for many-objective evolutionary algorithms. arXiv:2404.12746 [cs.NE]

[16] 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.

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

[18] 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.

[19] Weijie Zheng and Benjamin Doerr. 2023. Theoretical Analyses of Multiobjective Evolutionary Algorithms on Multimodal Objectives. *Evolutionary Computation* 31 (2023), 337–373. https://doi.org/10.1162/evco_a_00328

[20] 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, 20874–20882.

[21] Weijie Zheng, Mingfeng Li, Renzhong Deng, and Benjamin Doerr. 2024. How to use the Metropolis algorithm for multi-objective optimization?. In *Conference on Artificial Intelligence, AAAI 2024*. AAAI Press, 20883–20891.

[22] 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.