Multi-objectiveness in the Single-objective Traveling Thief Problem

Mohamed El Yafrani
LRIT URAC 29, Faculty of Science
Mohammed V University in Rabat
m.elyafrani@gmail.com

Shelvin Chand
University of New South Wales
shelvin.chand@student.adfa.edu.au

Aneta Neumann
Optimisation and Logistics
The University of Adelaide
aneta.neumann@adelaide.edu.au

Belaïd Ahiod
LRIT URAC 29, Faculty of Science
Mohammed V University in Rabat
ahiod@fsr.ac.ma

Markus Wagner
Optimisation and Logistics
The University of Adelaide
markus.wagner@adelaide.edu.au

ABSTRACT

Multi-component problems are optimization problems that are composed of multiple interacting sub-problems. The motivation of this work is to investigate whether it can be better to consider multiple objectives when dealing with multiple interdependent components. Therefore, the Travelling Thief Problem (TTP), a relatively new benchmark problem, is investigated as a bi-objective problem. The results indicate that a multi-objective approach can produce solutions to the single-objective TTP variant while being competitive to current state-of-the-art solvers.

CCS CONCEPTS

•Computing methodologies → Heuristic function construction; Randomized search;

KEYWORDS

Interdependence; Multi-component problems; Evolutionary Multi-objective Optimization; Travelling Thief Problem

ACM Reference format:
Mohamed El Yafrani, Shelvin Chand, Aneta Neumann, Belaïd Ahiod, and Markus Wagner. 2017. Multi-objectiveness in the Single-objective Traveling Thief Problem. In *Proceedings of ACM Genetic and Evolutionary Computation Conference, Berlin, Germany, July 2017 (GECCO'17),* 2 pages.
DOI: 10.475/123_4

1 MOTIVATION

In 2013, Bonyadi et al. proposed a benchmark problem called the Travelling Thief Problem (TTP) [1, 8]. TTP is a combination of the Travelling Salesman Problem (TSP) and the Knapsack Problem (KP). The goal of proposing such a fictional problem was to provide a more realistic academic model that simulates the research on interdependence of sub-problem in multi-component problems. In the original paper, the authors proposed two versions of the problem: a mono-objective TTP (TTP1) and a bi-objective TTP (TTP2). However, almost all published papers we are aware of are investigating the mono-objective version.

The problem supposes that a thief with his rented knapsack is willing to visit a set of cities. Each city contains a number of items, each having a value and a weight. The general goal of the problem is to help the thief find a path and picking plan in order to maximize the gain, minimize the lost, or find a trade-off between objectives depending on the problem's model.

In TTP's terminology, two main functions are defined: The first is **the travelling time**, which corresponds to the time taken by the thief to visit all the cities, pick up some items, and get back to the initial city. The second is **the profit**, which represents the total value of all stolen items. Additionally, in order to achieve the interdependence, two conditions are also proposed: **variable speed**, a constraint which supposes that the speed of the thief depends on the knapsack load; and **value drop**, a constraint which implies that the value of an item decreases with time.

Depending on how these functions and conditions are combined, different versions of the TTP can be created. It is clear that a TTP model can be formulated as a single objective problem or a bi-objective problem. We believe that the TTP1 formula is a simple scalarization of a multi-objective problem by nature—and therefore a multi-objective approach that does not consider the scalarization should be able to produce solutions covering a wide range of interdependencies.

The motivation of this work is to investigate multi-component problems as multi-objective ones by taking the TTP as a benchmark problem. Therefore, we are investigating the TTP as a bi-objective problem by considering traveling time and profit as the overall objectives.

Our investigations of this bi-objective model show that the best known TTP solutions can be found in the Pareto set region produced by our EMOA. It is even able to compete with three of the best algorithms for the TTP and find better solutions for the single objective model implicitly. For decision makers in the real-world who encounter multi-component problems, this can mean that comparable or even better solutions can be achieved if a multi-objective approach treats the different components as equally important.

2 PROPOSED APPROACH

Our proposed algorithm is built around the NSGA-II [3] framework as implemented in jMetal [4]. Instead of specifying the stopping

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'17, Berlin, Germany
© 2016 Copyright held by the owner/author(s). 123-4567-24-567/08/06…$15.00
DOI: 10.475/123_4

criteria in terms of function evaluations or generations, we use the 10 minute stopping criteria commonly used in the TTP literature. We define two disruptive operators and two local search heuristics as NSGA-II mutations. In the absence of an effective crossover operator, we use the *Null Crossover* to simply clone selected solutions. At the end of each generation the solutions are sorted based on dominance (non-dominated sorting operator). Solutions in each front are further sorted based on their crowding distance. Based on these two operators, the solutions for the next generation are selected. We will refer to this algorithm as *EMOA-TTP*.

3 EXPERIMENTS & RESULTS

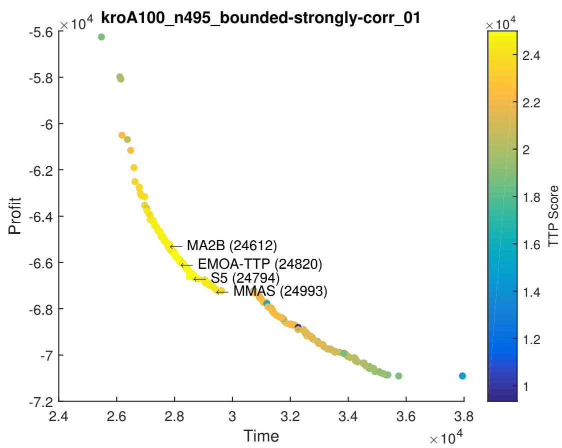

Figure 1: The obtained Pareto front for one TTP instance. The colors represent the TTP score.

As a first representative example, we show in Figure 1 the obtained Pareto front for one instance. The figure shows that *EMOA-TTP* was able to obtain a set of solutions that represent a trade-off between time and profit. We can also see that the best solutions regarding the TTP score are concentrated in the knee region of the Pareto front. In addition, the solutions obtained using *MMAS* [9], *MA2B* [5], and *S5* [6] are always close to the knee region. This shows that the single objective model is contained in the bi-objective model we investigated, as the TTP score is a simple scalarization of a bi-objective problem by nature.

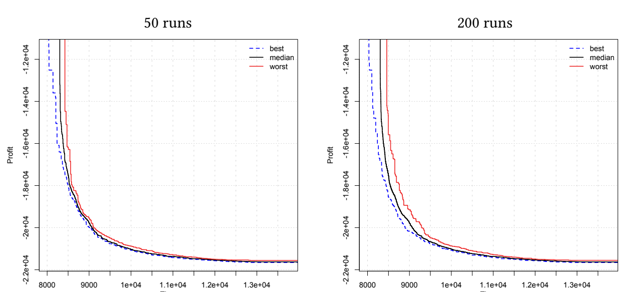

Figure 2: Attainment surface plots for 50 and 200 (from left) independent runs of the EMOA-TTP for the instance *berlin52_n255_uncorr-similar-weights_01*.

While the original single-objective TTP formulation allows for a straightforward comparison of objective scores, interpretations of the bi-objective results are more complex. We use the Empirical Attainment Function (EAF) [2, 7] to provide a graphical interpretation of the quality of the outcomes. In Figure 2 we show the variance of solution distribution depending on the number of runs of *EMOA-TTP*. The algorithm is run 50 and 200 times (from left to right) with randomly chosen seeds. Within each plot, the best results *EMOA-TTP* are shown by the best attainment surface (dashed blue line) and the worst results are shown by the worst attainment surface (solid red line). The median attainment surface corresponds to the region attained by 50% of the runs, which allows us to examine the median case quality of the attained objective vectors.

As we can see, the locations of the attainment surfaces are changing with increasing numbers of independent runs; in particular, the area between the worst and best surfaces increases, as expected for a stochastic algorithm. This means that the distribution of solutions varies and that multiple runs of *EMOA-TTP* have to be performed to compare a statistically significant differences between them at a later stage.

Our future work will include investigations to explore the benefits of a multi-objective approach as a single-objective solver, and as a contributor to TTP algorithm portfolios [10].

ACKNOWLEDGMENTS

This research has been mainly supported by an IEEE Computational Intelligence Society Graduate Student Research Grant and by an ARC Discovery Early Career Researcher Award DE160100850. The experiments were performed using supercomputing resources provided by the Phoenix HPC service at the University of Adelaide, Australia.

REFERENCES

[1] Mohammad Reza Bonyadi, Zbigniew Michalewicz, and Luigi Barone. 2013. The travelling thief problem: the first step in the transition from theoretical problems to realistic problems. In *Evolutionary Computation (CEC), 2013 IEEE Congress on*. IEEE, 1037–1044.

[2] Viviane Grunert da Fonseca, Carlos M. Fonseca, and Andreia O. Hall. 2001. Inferential Performance Assessment of Stochastic Optimisers and the Attainment Function. In *Evolutionary Multi-Criterion Optimization, First International Conference, EMO 2001, Zurich, Switzerland, March 7-9, 2001, Proceedings*. 213–225.

[3] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and Tamt Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.

[4] Juan J Durillo and Antonio J Nebro. 2011. jMetal: A Java framework for multi-objective optimization. *Advances in Engineering Software* 42, 10 (2011), 760–771.

[5] Mohamed El Yafrani and Belaïd Ahiod. 2016. Population-based vs. Single-solution Heuristics for the Travelling Thief Problem. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 (GECCO'16)*. ACM, 317–324.

[6] Hayden Faulkner, Sergey Polyakovskiy, Tom Schultz, and Markus Wagner. 2015. Approximate Approaches to the Traveling Thief Problem. In *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference*. ACM, 385–392.

[7] Manuel López-Ibáñez, Luís Paquete, and Thomas Stützle. 2010. Exploratory Analysis of Stochastic Local Search Algorithms in Biobjective Optimization. In *Experimental Methods for the Analysis of Optimization Algorithms*, Thomas Bartz-Beielstein, Marco Chiarandini, Luís Paquete, and Mike Preuss (Eds.). Springer, Berlin, Germany, 209–222. DOI: http://dx.doi.org/10.1007/978-3-642-02538-9_9

[8] Sergey Polyakovskiy, Mohammad Reza Bonyadi, Markus Wagner, Zbigniew Michalewicz, and Frank Neumann. 2014. A comprehensive benchmark set and heuristics for the traveling thief problem. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*. ACM, 477–484.

[9] Markus Wagner. 2016. Stealing Items More Efficiently with Ants: A Swarm Intelligence Approach to the Travelling Thief Problem. In *Proceedings of the 10th International Conference on Swarm Intelligence, ANTS 2016*. Springer International Publishing, Brussels, Belgium, 273–281.

[10] Markus Wagner, Marius Lindauer, Mustafa Mısır, Samadhi Nallaperuma, and Frank Hutter. 2017. A case study of algorithm selection for the traveling thief problem. *Journal of Heuristics* (2017), 1–26.