



## 1 Motivation

Today, containers are widely applied as means to host applications in an elastic way. Despite this, existing approaches to optimize the allocation of containers on Virtual Machines (VMs) are rather simplistic, or do not cover the whole spectrum of computational resources [1, 2]. For instance, privacy aspects or hybrid clouds are so far not regarded to a large extent [3].

The goal of this thesis is to extend the state-of-the-art in resource allocation for containers by conceptualizing and implementing an optimization model which takes into account user requirements and places containers accordingly on VMs. In addition, the size/type of the VMs is also determined by the outcome of the optimization. Since the underlying optimization problem is NP-hard, the student also has to come up with heuristics, e.g., based on Genetic Algorithms, to solve the optimization problem.

## 2 Work Description

- Literature work on resource optimization, containers, and cloud computing. Especially, the student needs to identify open gaps in resource optimization for containers by conducting a systematic literature review.
- Modeling of an optimization problem and according heuristics.
- Implementation of the model and the heuristics.
- Quantitative evaluation.

## 3 Further Information

**Start:** Immediately (might also be later)

**Basic Requirements:** Very good implementation skills; knowledge in optimization is a big plus

## References

- [1] Y. Al-Dhuraibi, F. Paraiso, N. Djarallah, and P. Merle. Elasticity in Cloud Computing: State of the Art and Research Challenges. volume 11, pages 430–447. 2018.
- [2] R. Buyya, S. N. Srirama, G. Casale, R. N. Calheiros, Y. Simmhan, B. Varghese, E. Gelenbe, B. Javadi, L. M. Vaquero, M. A. S. Netto, A. N. Toosi, M. A. Rodriguez, I. M. Llorente, S. D. C. di Vimercati, P. Samarati, D. S. Milojevic, C. A. Varela, R. Bahsoon, M. D. de Assunção, O. Rana, W. Zhou, H. Jin, W. Gentzsch, A. Y. Zomaya, and H. Shen. A manifesto for future generation cloud computing: Research directions for the next decade. *ACM Computing Surveys*, 51(5):105:1–105:38, 2019.
- [3] P. Hoenisch, I. Weber, S. Schulte, L. Zhu, and A. Fekete. Four-fold Auto-Scaling on a Contemporary Deployment Platform using Docker Containers. In *13th International Conference on Service Oriented Computing*, volume 9435 of *Lecture Notes in Computer Science*, pages 316–323. Springer, Heidelberg Berlin, 2015.

-3 -2 -1 0 1 2 3

Conceptual (Analytical)

-3 -2 -1 0 1 2 3

Empirical (Simulation)

-3 -2 -1 0 1 2 3

Practical (Implementation)

-3 -2 -1 0 1 2 3

Literature Work