

# 論文の内容の要旨

## Applied and Scalable Optimization of Long-term and Network-aware Service Compositions

(ネットワークを考慮した長期間に渡るサービス  
合成におけるスケーラブルな最適化の適用)

氏名 クライン・アドリアン・ヘルムート・ダビド

Computing optimal service compositions in terms of their non-functional, Quality of Service (QoS), properties, is one of the key problems in realizing the vision of Service-Oriented Computing (SOC), and is known as the NP-hard QSC problem. The spread of Service-Oriented Computing and the rapid increase in the number of available services have spurred various research investigations over the last 10 years. Once software components are defined as services, they can be easily combined to achieve more complex functionality, even by less technical users. Services provide loosely coupled access over the network through clearly defined APIs. This enables robust coupling of software components across both network and company borders. It also makes the SOC paradigm a natural fit for the recent trend of moving an increasing number of software components into the cloud.

This thesis investigates the problem of effectively and efficiently computing near-optimal service compositions for long-term and network-aware settings. By extending the QSC problem for these settings, we close the increasing gap between the standard formalization and the current reality of service compositions. As both the problem's complexity and its search space increase for these extensions, we also build and improve upon ongoing research to achieve a good scalability by using domain knowledge to guide our optimization for the QSC problem.

As service compositions become a crucial part of the software landscape, compositions are often carefully specified and then used over an extended period of time. In such long-term settings, optimizing for multiple executions at once allows our approach to provide significant QoS benefits to users, especially, when users specify tight QoS constraints such as a monthly budget or a high minimum reliability. For this purpose, we compute probabilistic and time-dependent execution policies through Linear Programming and a custom Genetic Algorithm (GA), which extend the standard static mapping between the parts of a composition and concrete services. Through an adaptive GA we achieve a reasonable scalability for the latter even though the complexity of the QSC problem increases due to the number of time-dependent choices.

With the growing distribution of services within compositions, e.g. in the context of clouds, the influence of the network on the overall QoS of compositions keeps increasing. We propose both a network-aware modeling of QoS and a network-aware optimization. This includes a formalized distributed architecture, a generic network model and a customized self-adaptive GA. Modeling QoS in a network-aware manner requires distinguishing even between different physical instances of the same service by the same provider, as users will have different experiences depending on the QoS of the network between them and these instances. Thus, the number of optimization choices increases. Furthermore, considering the network also introduces strong dependencies between connected services within compositions. Therefore, the complexity of the problem increases significantly and exploring the search space in an uninformed way becomes less effective. In order to solve this challenge, our optimization approach is aware of the network and its custom GA operators make informed probabilistic decisions by using domain knowledge about the network. As other QoS properties unrelated to the network, such as price or reliability, still have to be optimized as well, this supports our choice of a GA which allows our approach to seamlessly use both domain-specific and general operators at the same time. Our custom self-adaptation rules assure that the most appropriate operators are chosen depending both on the concrete problem instance and the user's QoS preferences, thus, guaranteeing the generality of our approach for the QSC problem. We evaluate our approach based on an extensive (externally provided) network dataset against standard GAs which represent state-of-the-art algorithms for the QSC problem and against a Dijkstra algorithm exclusively optimizing the network latency. In terms of network latency, our approach consistently achieves a good approximation ratio for all evaluated problem sizes compared to the optimal latency computed by Dijkstra, while the approximation ratio of the standard GAs deteriorates with increasing problem size. In terms of other QoS, it is on par with or slightly better than standard GA approaches. The scalability of our approach beats standard GAs, and in case of an increasing number of services instances, Dijkstra is outperformed by several orders of magnitude.