論文の内容の要旨

Robust Workflows by Applying Functional Clustering on Multi-Objective Service Composition

（多目的のサービス合成における，
機能クラスタリングの適用による
ロバストなワークフローの構築）

氏名　　ワーグナー　フロリアン　ステファン

Service-Oriented Computing is a widely applied paradigm for establishing loosely coupled components on top of company networks and the World Wide Web. Services are reusable and interoperable entities that encapsulate business functionality behind standardized interfaces. Apart from the functional interface descriptions, service providers declare the Quality of Service (QoS), for instance the price and response time, of their services. By composing services, a virtually unlimited number of functionalities can be provided.

In the past decade, two main approaches towards service composition have emerged: service planning and service selection. The former approach employs AI planning to automatically generate workflows, whereas the latter refines given workflow templates. We claim that both approaches are insufficient to align the functional and non-functional requirements of the user. Planning lacks scalability and moreover neglects QoS aspects. Service selection on the other hand is inflexible; in case the required functionality changes, the template might become unusable. Moreover, functionally diverse services are not considered in selection algorithms. Both approaches rely on expensive ad-hoc replanning at runtime to deal with service failures. The major problem with such sequential application of composition and replanning is that it ignores the potential costs of replanning during the initial composition and they consequently are hidden from the user.

In this thesis, a holistic approach towards service composition is proposed that integrates service planning and service selection, bridging the gap between those two research areas. More specifically, we present the three following main contributions.

First, we propose a functional clustering of services that helps to determine backup services beforehand and to consider service functionalities. For each cluster, representative services are computed that are used in the following two contributions. An extended QoS model is introduced that estimates the impact of service failures on the QoS of a workflow.

Second, we present our planning algorithm, Keikaku, that computes workflows on an abstract level. By considering aforementioned representatives instead of single services, unnecessary comparisons of services can be avoided. This addresses the scalability issues of planning tools.

Third, we extend the original problem definition of service selection by considering functionally diverse services. In this setting, local optima are more likely, which leads to poorer results of heuristic algorithms that explore the search space in a more or less random manner. We propose two domain-specific extensions of Genetic Algorithms, called Shuuzen and Shuuri, that leverage our functional clustering as background knowledge to compute solutions efficiently.

We further present a holistic framework that unifies these three contributions and present a survey on the current usage of service computing.