1 Introduction

Fostered by the Internet of Things (IoT), the Internet of Services (IoS) has become the “new economy” in the Internet. Combining services from independent providers enables other providers to offer value-added, composite services to their customers.

This thesis contributes to the creation, selection, and provision of services by harvesting quality-related information from reviews and ratings for services. The central theme is the support of providers (of both services and composite services) and customers in deciding for the best available services. Therefore, reviews and ratings for composite services are *decomposed* into ratings for the individual services, which form the composite services.

Section 1.1 describes the setting, in which the contributions of this thesis are made. Section 1.2 introduces the individual contribution made, while section 1.3 gives an overview on already published parts of the thesis. Finally, the structure of this thesis is explicated in section 1.4.

1.1 Motivation and Setting

Service provision in the IoS concerns several stakeholders, which are presented in the following. Including reviews and ratings as additional sources of quality information leads to a four-phase IoS service provision process as depicted in figure 1.1.

	Phase 1: Creation	Phase 2: Orchestration	Phase 3: Provision	Phase 4: Use
Stakeholder	Component service providers	Orchestrators	Marketplaces	Customers
Activity	produce and offer	combine* component services	offer composite services	use and rate
Target	their services	to composite services	to customers	composite services

*component service providers may refuse unfair composition

Figure 1.1: Four-phase service provision process, including reviews and ratings.

1. During the first phase, self-contained services (**component services**) are created, i.e., designed, implemented, and tested. After the creation, independent service providers, the **component providers**, offer their services to potential users.

2. In the second phase, **orchestrators** use existing component services to create **composite services**. Thereby, the existing component services are combined and merged together with technologies from the field of service oriented architectures. On a nonfunctional level, aspects of fairness can be considered during composition.

3. The third phase is the composite service provision phase. Orchestrators offer their composite services in **marketplaces** to potential **customers**.

4. Customers use composite services they find in marketplaces. Thereby, the composite service is executed by executing the respective component services in the order defined by the orchestrator. After the last

component service has finished, the composite service result is delivered to the customer. Based on the quality of the result, customers can leave **feedback** for the composite service.

1.1.1 Stakeholders

In the setting of this thesis, the aforementioned stakeholders are interconnected as shown in figure 1.2.

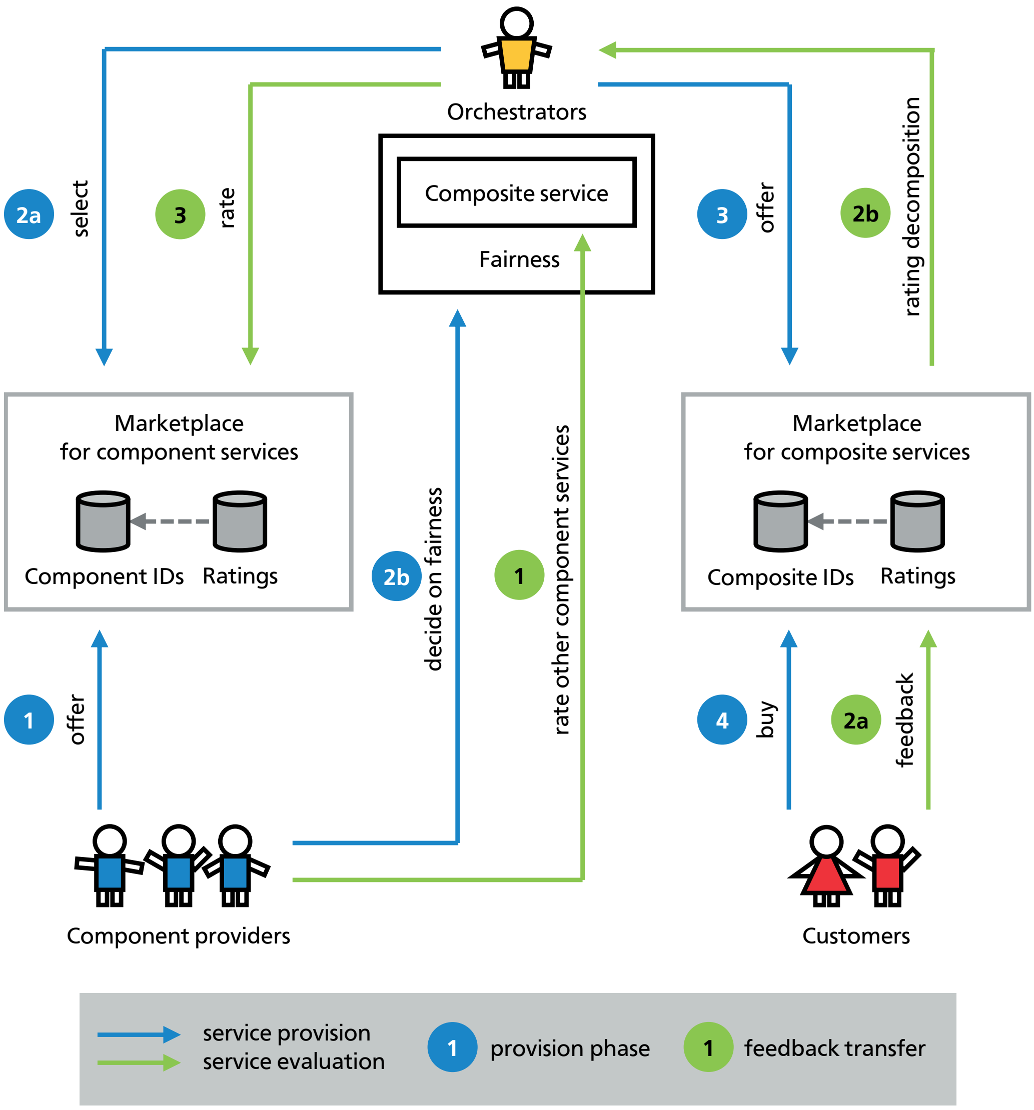

Figure 1.2: Relations and connections in between the stakeholders of the composite service setting.

Component Provider

Component providers create and offer component services, which are the base of composite services. The customers of component providers are orchestrators that use the offered component services in order to create composite services. Using, i.e., executing, component services from within a composite service is to be payed by the orchestrator. The component provider profits from these payments.

Component providers are a source of feedback, when their component services are used within a composite service (cf. **inter-component rating** and **performance functions** in chapter 4).

Orchestrator

Orchestrators combine component services to composite services. These composite services are offered to customers via marketplaces. Whenever a customer uses a composite service, the component services that form the composite service are executed in the predefined order, creating a value chain in which results of component services become input of other component services. The endmost component service result is the output of the whole composite service and given to the customer. Customers pay to orchestrators for using composite services.

By decomposing feedback from customers (cf. chapter 6), orchestrators obtain ratings for component services.

Aspects of distributive fairness can be considered during the selection of potential component services. A formal model for distributive fairness, which can be used to determine the fairness of revenue distribution among component providers, is presented in chapter 5.

Customer

Customers use or "consume" composite services found in marketplaces and offered by orchestrators.

As receivers of the composite service output, customers a source of feedback. Customer feedback usually applied to composite services a whole and is subject to uncertainty. Chapter 6 addresses customer feedback and proposes ways to derive ratings for individual component services within composite services.

Marketplace

Marketplaces are SOA infrastructures, where providers offer their services to potential customers. Customers can retrieve a list of composite services that fulfill their functional requirements. The same holds for orchestrators, who find component services in marketplaces.

In order to fulfill their function, marketplaces make component services and composite services recognizable by digital identifiers. Authentic digital identifiers enable orchestrators and customers to distinguish different services and to assign feedback. The contribution presented in chapter 3 uses **computational trust** (cf. section 2.1) to improve the security of digital identifiers based on certificates.

Besides brokering services from providers to users (either from component providers to orchestrators or from orchestrators to customers), marketplaces are also a place, where feedback can be collected and retrieved. Customers submit their feedback on composite services to marketplaces in order to aid other customers in their purchase decisions (for communicating decomposed feedback to human customers, refer to chapter 7) and, furthermore, to help orchestrators in improving their composite services. Orchestrators use feedback to select component services by nonfunctional aspects.

The boundary between component services and composite services is blurry: component services might themselves act as component services within other composite services. Thus, the two marketplaces shown in figure 1.2 might as well be one large marketplace, in which services are offered to orchestrators and customers as well.

1.1.2 Trust Relationships

The IoS setting, as introduced in the previous section 1.1.1 and figure 1.2 (page 21), includes five trust relationships relevant for the contributions of this thesis.

1. **Trust from orchestrators to component providers.** Orchestrators build trust relationships to component providers, trusting in their ability to provide component services with the advertised quality. Ratings for component services are used in chapter 4 as data source for calculating this trust.

2. **Trust from component providers to orchestrators.** When component services are composed to composite services, the component providers trust the orchestrators to attribute the contribution of each component service in a fair manner. A unified model for distributive fairness is contributed in chapter 5.

3. **Trust from customers to component providers.** Customers of composite services trust the providers of the used component services to provide their services with the quality advertised by the composite service. Two approaches to decompose reviews and ratings for composite services into ratings for individual component services are contributed in chapter 6: a decomposition-supporting review system that supports customers in giving decomposed ratings and an automated mechanisms to derive component service ratings from a composite service rating.

4. **Trust from orchestrators and component providers to the marketplaces.** Reviews and ratings need to be assigned to recognizable entities in order to available for later trust calculation. In chapter 3, the binary trust model of current certificate infrastructures is evolved in order to improve the authenticity of digital identifiers.

5. **Trust from customers and orchestrators to the marketplace.** Having ratings for component services available supports customers in choosing which composite service to use and orchestrators in selecting appropriate component services for their composite services. Thereby, trust in a composite service is multicriterial information consisting of trust information for each component service. T-Viz, a visual multicriterial trust visualization is contributed in chapter 7 and is designed to support customers and orchestrators with an intuitive and easy to comprehend graphical trust representation.

1.1.3 Example Service

In order to illustrate the mechanisms proposed in this thesis, the running example of an illustrated book service is used. This section presents the illustrated book service in detail.

The task of the example service is to create illustrated books for various topics, for example, the beaches of Southern France. The result of the service is a printed book, which combines color photographs with matching pieces of text in an appealing layout. Customers of illustrated books are generally interested in obtaining a book that offers unusual insights into a topic in form of illustrations side by side with well-researched background information. Moreover, customers expect high quality print as illustrated books commonly serve as gifts.

Figure 1.3 visualizes the components of the illustrated book service and their interrelations. The service is composed of three component services, namely the *Photo & Text component*, the *DTP Layout component*, and the *Printing component*. The customer gives a topic and the expected number of pages to the orchestrator, which then instructs the photo & text component accordingly. The three component services are used in the following sequence:

1. The **Photo & Text component** supplies photos and matching pieces of textual background information.

2. The **DTP Layout component** arranges the supplied photos and text in a suitable and appealing layout, ready for printing.

3. Finally, the **Printing component** is responsible for printing and binding the book.

However, this composite nature is hidden from the customer, as the internal structures of the component services are usually unknown to the orchestrator.

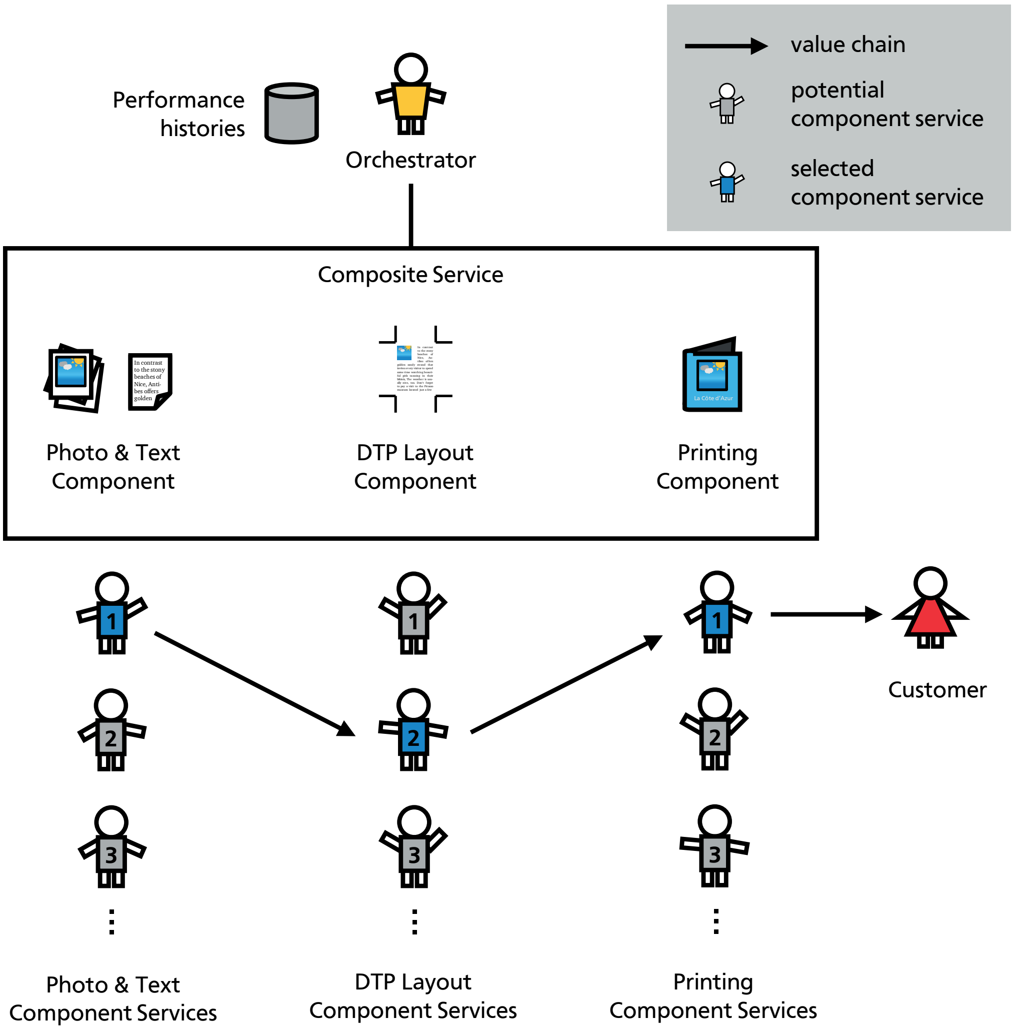

Figure 1.3: The example illustrated book service.

Service Quality

For every component service, the orchestrator can choose from a set of services that all serve the intended purpose and only differ in quality.

The photo & text services deliver photos of different esthetics and image quality as well as differently well researched background information. Differences in the quality of the layout, in terms of optical clarity and typesetting of textual background information, distinguish the DTP layout services. The printing services have different qualities of workmanship regarding the print output itself, but also regarding the materials and durability of the book cover.

It is to note that quality and price are related. Often, higher quality also implies higher prices and orchestrators will usually try to find an optimal relation between overall quality of the combined services and prices. Thereby, it is possible to offer the same composite service in different qualities and prices for distinct groups of customers.

Component Service Selection

For all component services known to the orchestrator, i.e., those component services that contributed at least once to the illustrated book service, the orchestrator stores **performance histories**. These performance histories indicate the quality to expect from the respective component services based on their historical output quality. Everytime the illustrated book service is used, the orchestrator decides which component services to use from the set of available component services. It is the mere interest of the orchestrator to choose those component services which are expected to perform best in order to obtain the best possible output quality for the whole composite illustrated book service.

Composite Structure

According to Medjahed's and Bouguettaya's classification of composite services [MB05] (later extended in [NMB09]), the illustrated book service is a *horizontal* composite. This means, the components the service is composed of depend on each other in order to create the composite service output. The output of the first component service becomes the input for the second one et cetera. The output of the last component service in this chain is the overall output of the composite service.

Other composite structures are introduced in section 2.3.1.

1.2 Challenges, Contributions and Evaluation Methods

This thesis addresses the interface between customers and providers of IoS composite services. The contributions target feedback pre-processing as well as decision-making based on feedback for both customers and service providers.

In this section, the fields of contribution are introduced. Afterward, the specific contributions are discussed. For every contribution, the evaluation method is presented and justified. Additionally, central results are presented.

1. Section 1.2.1 introduces the field of digital identities implemented by certificates and hierarchical certificate infrastructures. The related contribution, which improves the authenticity of digital identities, is found in chapter 3.

2. Section 1.2.2 introduces the field of rating composite services by deriving ratings for component services from external information. Chapter 4 presents the contribution of inter-component rating and performance functions, which enable orchestrators to profit from quality information supplied by component providers.

3. Section 1.2.3 introduces the field of distributive fairness, i.e., the evaluation how fair sharing a good among entities is. In chapter 5, a unified fairness model (along with six types of fairness, called *fairness classes*) is contributed. This model can be used by component providers to reason on the fairness of their payment in a composite service.

4. Section 1.2.4 introduces the field of customer-generated feedback, as in, e.g., reviews and ratings. Chapter 6 proposes a decomposition-supporting review system, which supports customers in giving decomposed feedback. Moreover, chapter 6 contributes an automated approach to decompose ratings given by customers for a composite service into ratings for the component services used in the composite service.

5. Section 1.2.5 introduces the field of visually communicating trust statements to human users. Chapter 7 contributes T-Viz, a trust visu-

alization that concurrently shows multiple trust values, for example, for all component services of a composite service.

1.2.1 Trust in Digital Certificate Infrastructures

Identifying people and verifying their identities is usually easier in daily life than in digital settings: people have individual physical and behavioral cues that can be detected and recognized by other people. In the digital world, however, those cues might be unavailable. People or any other kind of entity, like, services and service providers, are to be identified and recognized by their **digital identities**. In this thesis, digital identities are used to assign feedback to entities: orchestrators use digital identities to (re-)identify component services and keep track of their historical performances. Customer feedback is also bound to the digital identities of IoS composite services.

In comparison with real world identities, digital identities have distinct properties, negatively affecting the verification of authenticity:

- A single entity can hold multiple digital identities concurrently; several entities can share one digital identity.
- Digital identities can be maliciously taken over by other entities without this becoming noticeable.
- Digital identities can be copied.
- Some common implementations of digital identities require identities to expire. In this case, a digital identity must be renewed and changes during this process.

A common way to implement digital identities are **digital certificates**, which bind a name to a cryptographic public key. Such a digital certificate is signed by a trusted **certification authority** (CA). Only the owner of a certificate is supposed to know the private key belonging to the public key. Therefore, only the owner can prove the name in the certificate is his or her digital identity. All of this only holds under the assumption that the CA is trustworthy, i.e., the CA only signs valid digital certificates.

Application of Computational Trust to Hierarchically Structured Certificate Infrastructures

In chapter 3, the contribution to digital identifiers is presented and evaluated: by applying methods and models from computational trust (see section 2.1 for an introduction to computational trust) to hierarchically structured certificate infrastructures, the security of digital identifiers is increased. CAs are organized in a hierachical structure, which means that CAs can sign digital certificates of other CAs. Thus, the user must not know all trustworthy CAs, but only a set of **root CAs**. Starting from those root CAs, digital certificates can be evaluated by following a chain of digital certificates, as shown in figure 1.4.

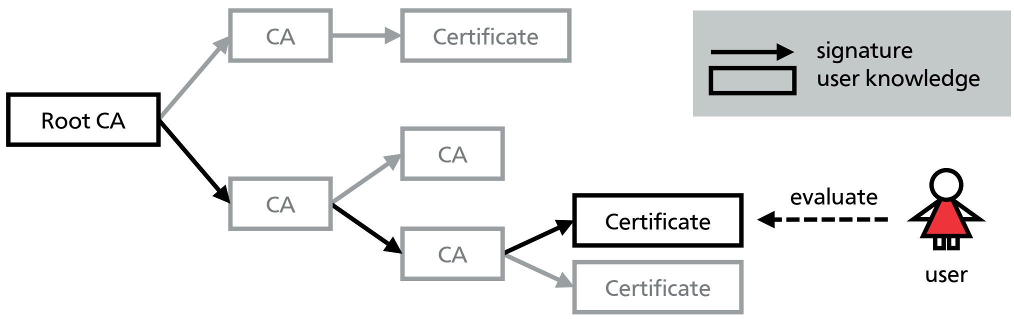

Figure 1.4: Certificate chain in a hierarchically structured certificate infrastructure.

The presented application of computational trust realizes local **trust views** [BVBM13, BVC+14]. Instead of relying on a vendor-defined list of 100% trustworthy root CAs (deployed with operating systems or web browsers), users can use trust views to "learn" individual trust values for the CAs they actually need. Thus, the attack surface is effectively reduced to only those CAs the respective user needs to trust. Moreover, by sharing their trust view in reputation systems as proposed in [BVC+14, CBV+], users profit from the knowledge of similarly behaving users.

Evaluation

In order to evaluate the effectiveness of the approach presented in chapter 3, the impact on the size of attack surface, that is, the number of trusted CAs, has been tested on the web public key infrastructure (Web PKI), which is the largest hierarchical infrastructure of digital identities. The concept of trust views with computational trust has been implemented and evaluated on browsing histories obtained from real web users. Several of the browsing histories used for the evaluation span over many years and thus, contain every digital certificate a user is confronted with over a long period of time.

The evaluation (section 3.2.1) shows that the attack surface is reduced to those CAs which the user needs to verify the identities of the web hosts it encounters. Even though reputation systems increase the attack surface slightly, the reconfirmation rate for newly encountered certificates is reduced. Over all, improving the security model of the Web PKI with an advanced trust model increases the ability of users to authenticate web host identities.

1.2.2 Inter-Component Rating and Performance Functions

Authentic digital identifiers make component services and composite services recognizable in repeated use. However, orchestrators compose composite services only with limited information on the performance of the selected components. Customer feedback usually treats the output of a composite service as a monolithic unit. This holds especially for horizontally composed services, in which intermediate results might be invisible for customers.

In order to offer a composite service with the highest possible output quality, the orchestrator rather wants to obtain feedback on components than feedback on the whole composite service. With feedback regarding individual components, orchestrators can select and combine the best known component services, or replace component services that negatively affect the quality of the composite service.

Inter-Component Rating and Performance Functions

This thesis contributes **inter-component rating** (ICR) and **performance functions** in order to address obtaining feedback for individual components of a horizontal composite service in chapter 4. A combination of ICR and performance functions enables orchestrators to collect ratings for all component services individually. Thus, issues that impact the overall output quality of a composite service can be localized and attributed to specific component services.

- **ICR** incorporates the providers of component services as a source of feedback [Vol13]. By rating the quality of the input a component service receives, feedback on the output quality of the preceding component service becomes available to the orchestrator. Moreover, component providers are experts on evaluating their input. Thus, honest ratings from component providers should be much more accurate than customer feedback.

- **Performance functions** model the relation between the input quality of a component service quality and its output quality. Thus, performance functions define the **ability to perform** under specified constraints [VSM14]. This is necessary as the output quality of a component service in a horizontal composite depends on at least two factors: (i) the component service's performance and (ii) the quality of the input. Even the best component service will not be able to create output of high quality if the input quality is too low. By using performance functions, which assign every input quality the respective output quality the component service will deliver, component services are protected from false blame resulting from low-quality input.

With the information gained from applying ICR and performance functions, the orchestrator can initiate appropriate measures to improve the overall output quality, for example, by replacing component services with better rated ones. Moreover, the obtained ratings can be used by orchestrators for a second purpose: tracking performance histories of component services. With information on the past performance of component services, their future performance can be predicted. Hence, orchestrators can make better decisions when selecting component services by choosing those with the best predicted performance.

Evaluation

In order to evaluate the effect of ICR and performance functions, an extensive simulation study was implemented. Therein, the output qualities of two composite services were compared over 100 rounds. The results are average values from over 1.000 simulation runs. The first composite service was provided by an orchestrator that has only access to overall ratings, i.e., ratings for the composite's output quality. The second operator applied ICR and performance functions. Both operators chose from the same set of component services and tracked the performance of component services by the means described above.

Moreover, the time taken to notice a change in the performance of a component service was measured in the same scenario over 500 rounds. Even though the proposed combination of ICR and performance functions does not include fraud detection, the influence of fraudulent component providers was investigated. Fraudulent providers report better performance than their offered component services achieve while simultaneously trying to hide this fact from the orchestrator. For example, such a provider might report an untruthful low rating for the input given to its component service in order to cover its incapability to perform as good as the actual input quality would imply.

The simulation results show that an orchestrator that applies ICR and performance functions achieves higher composite performance (0.993) than an orchestrator without access to ratings for intermediate results (0.931). Moreover, the standard deviation of the qualities achieved by the first orchestrator is lower than for the second orchestrator. When selecting component services, applying ICR and performance functions enables orchestrator to better distinguish low-performing, regular-performing, and high-performing component services.