Assessing Web Services Quality for Call-by-Call Outsourcing

Wolf-Tilo Balke
Computer Science Department
University of California
Berkeley, CA, USA
balke@eecs.berkeley.edu

Atta Badii
Department of Computing
University College Northampton
United Kingdom
atta.badii@northampton.ac.uk

Abstract

The process of decision making regarding the outsourcing of various tasks in an enterprise's value chain involves consideration of a number of important economic trade offs. Outsourcing decisions in the past have typically been of a strategic nature involving long-term commitment. However, given the emerging technologies of Web services and the growing variety of service providers it seems feasible to utilize the Internet's ubiquity and pervasiveness for outsourcing at least more simple tasks dynamically in a call-by-call manner. However, to make the required decisions in each instance without having to manually assess the quality of each offer an automatic quality assessment based on a mutual service level agreement is essential. In this paper we will investigate the necessary mechanisms and demonstrate their feasibility in a real world use case scenario.

Keywords: web services, quality of service, service level agreement, utility assessment, outsourcing.

1. Introduction

The World Wide Web has gradually made its way from a simple communication space to a fully-fledged enterprise space far beyond providing a space for mere web presence and advertising for major companies. The service requirements of the future dynamic and mobile commerce need to cater for mobile users and roaming executives in the grid-aware economy. They will essentially require an integrated service provisioning environment facilitating support for service specification, service level agreement composition, and monitoring including outsourcing relationship management services as highlighted by our previous research focusing on outsourcing relationship management and information symmetry for mutual benefit [2]. With the emergence of the Web services paradigm and the emerging transition to grid-based service oriented architecture, the idea of service and content provisioning over the Internet using its pervasiveness and ubiquity has become feasible for the first time and its impact on toady's economical structure is still uncertain.

The failure of previous approaches towards E-business must call for a re-examination of their readiness in particular their economic and cultural fit. It can be seen that they largely amounted to a poor fit missing some of the typical characteristics and common requirements of today's economic culture, e.g. provisioning of advocacy-trust services or user-centered information provisioning plus knowledge management services that allow individuals to efficiently cope with their everyday selection of services to support their daily organisational procurement, work and life styles see e.g. [22]. In this context key challenges remain to be overcome; for example the integration of discrete tasks within often dynamic and complex value chains, the decomposition and chaining strategies for complex tasks or workflow processes to enable competitive multi-sourcing, and, the personalisation of such services for B2B or B2C service provisioning. Since we cannot address all these issues within our paper we will limit our focus to cooperative service selection for outsourcing tasks and the necessary service level agreements.

Far from acting simply as a source of information services on demand, Web services should take on the well-suited additional role of becoming also an important platform for the outsourcing of certain procedures that fall outside an enterprise's core competencies. As such these procedures cannot be performed internally as a routine due to a lack of resources or expertise, although their execution is necessary in closure of certain value chains. However, the effective deployment of such call-by-call invocable services requires additional infrastructure components beyond the underlying protocols, inference languages or shared ontologies. Usually service provider or service portal manager can anticipate a wide variety of typical user interaction (stereotypes) and provide at least some basic 'ready-to-use' querying tools or domain ontologies. The provisioning of a platform for continuous monitoring of available service offerings, levels of service guarantees, automatic quality-of-service (QoS) assessments (i.e. quality of personalised service provisioning), estimates of execution costs and advanced personalisation in both finding and selecting the appropriate service for subsequent execution is vitally important for the take-up and sustainable success of web-enabled call-by-call outsourcing business.

In this paper we shall address some of these problems and propose preliminary approaches from our current research in provisioning of Web services in mobile user-centered environments. We shall begin here by first discussing an extended case study to highlight the implications for service provisioning models. We argue that with more effective techniques for assigning utilities based on semantic descriptions of services and vendors' offers, an automatic monitoring of applicable services can be made available to help establish a dynamic outsourcing paradigm even for more complex tasks. Thus outsourcing as a strategic business decision with necessary long-term commitments may at least for some tasks begin to give way to the more efficient paradigm of by *call-by-call outsourcing*.

This could evolve along similar lines to that of the personal telecommunications consumer market where exclusive long-term contracts with a single telephone company have been largely superceded by a call-by-call provisioning paradigm. Nowadays given time, place and distance, competing telephony services providers each offer their specific rate for each single phone call and thus the individual user is free to choose the cheapest provider in each instance. However, in this area of application the level of service agreement is quite clear: the provider is expected to supply an immediate telephone connection of 'sufficient' transmission quality and the price is fixed given the characteristics (from/to, duration, etc.) of the call. Generally users then can simply choose the provider offering the cheapest rate. In a consumer context going beyond the purchase of mere telephony services, negotiating an appropriate service level agreement for the provisioning may, however, prove considerably more demanding.

2. Strategic Planning and Outsourcing

For the following we will focus on simple flexible services that are particularly convenient for every day business life. This kind of services is often referred to as user-centered service, because usually it is offered directly to the user in a B2C fashion and unlike the fixed (and mostly well-understood) capabilities of B2B services, it involves a large amount of personalization to fulfil the expectations of a variety of end-users. We will show that when handling outsourcing decisions dynamically through monitoring, and optionally reacting to, interesting offers, business procedures can be essentially improved. Please note that we do not intend to generalize this approach to all outsourcing decisions. Of course, for dominantly strategic service procurement decisions, where longer term relationships, trust and confidence and above all service level sustainability issues are paramount (e.g. outsourcing a company's IT department), the outsourcing management will, for the foreseeable future, continue to include significant long-term outsourcing.

However, it is expected that once the operational efficiency of call-by-call outsourcing is established given adequate layers in the online transaction infrastructure, the frontiers of call-by-call outsourcing will push forwards gradually to include more and more of the outsourceable enterprise operations. This trend will be powerfully aided by the advent of new business models and new ways of working and living in a grid-aware global networked economy. It should also be noted that indirect advantages and differential incentives attendant with various outsourcing schemes like tax deductions, etc. will not be considered here. However it is clear that, given appropriate infrastructural support as highlighted above, for a large share of everyday smaller tasks as frequently deployed by most enterprises a web-assisted, *call-by-call outsourcing* paradigm can prove particularly useful and given the growing capabilities of Web services and the progress in the area of semantic Web, more and more of then complex tasks might become candidates for automatic call-by-call outsourcing.

2.1. A Use Case Scenario

Let us consider a company in which a large amount of staff travel has to be organized frequently, for instance a consulting business. In order to determine the best advice and solutions adapted for each customer, consultants have to work at the customer's premises from time to time, but also have to be present at their own headquarters for other projects or to achieve synchronization with the rest of the project team. It is quite common for companies which need to provide transportation for their employees on a regular basis to have a standing outsourcing agreement with e.g. a certain travel agent or a rental car company whereby, on demand, the rental car company would supply a car and would offer the company special rates -often in return for a guaranteed number of rentals per year. The mutual benefit obtained with such agreements is twofold:

- Having a certain amount of business guaranteed over a business cycle makes for improved economic planning for the rental car company thus enabling it to offer better terms and conditions of service to its customers; this can often mean charging specially lower rental rates overall for clients with whom it has such outsourcing agreements.
- Furthermore, the time and effort (for administration, advertising, etc.) spent by the rental company each time in securing contracts, i.e. getting (new) customers and on the other hand the administrative overhead incurred by the client company in re-acting anew to each and every instance of need for an appropriate rental company is minimized.

For instance if a consultant of a New York-based consulting company needs a car on a business trip to Boston, a nationwide rental car company might provide the service and there is no need for an outsourcer to phone a number of different Boston rental car companies to find a suitable provider and (manually) compare their prices, i.e. (administrative) tasks can be reduced and thus costs cut. Though cutting costs is a primary goal, on the other hand in each instance there may also be competing special offers by local providers which can be by far cheaper than the fixed price to be levied for such a journey per the outsourcing agreement and the outsourcer is usually bound by a contract to guarantee a minimum amount of business within a certain time span. Thus rather than cashing the gain in administrative costs cut, for a real outsourcing balance a customer would have to compare the pairs of (travel agents fee and fixed prices - administrative costs cut) versus (cheapest offer for each travel + administrative cost).

So from the outsourcer's viewpoint there are always advantages and disadvantages attendant with each of the two alternatives of either invoking or bypassing the existing rental car outsourcing agreement and the prize of reduced costs *overall* will belong to the client who is enabled to strike the best trade-off i.e. enabled to invoke call-by-call outsourcing thereby achieving a managed mix of alternative choices over a business cycle. Typically the relative benefits of each alternative can be readily deduced and in opting for the outsourced rental car as routine, the tradeoffs would include, for example:

+ special tariff applies; i.e. discount on each rental
+ cutting labour costs or relieving individual work loads
– long-term commitment, lack of flexibility
– having to guarantee a certain amount of business

From a business perspective especially the lack of flexibility is a major drawback. If we assume that at least in some instances our company can, without incurring outsourcing penalties, be freed from remaining a captive client so as to get better rates from another provider than the ones specified in the outsourcing agreement then this client would be able to minimise their overall costs, if the occasionally available better rates could be chosen in each *suitable* instance. Such *suitable* opportunities for exercising this freedom of choice can occur at some point in a given business cycle once the enterprise is able to ascertain that, over that cycle, it either has already fulfilled its minimum service purchase guarantees vis-à-vis the vendor or will assuredly be able to do so in the near future. However, finding the better rates to enable occasional bypassing of the outsourcing vendor and call-by-call purchasing of the rental service would need a monitoring of services which of course would increase labour costs, especially if after an assessment of competing rates none proved to be more attractive than those already offered by the existing outsourcing vendor. However, it is perfectly possible that the net gain from occasionally bypassing the regular provider will over time amount to significant savings if only such occasions could be efficiently spotted.

Thus with a long-term car rental service provisioning agreement already in place, the company would be ill-advised to devote resources also to the continuous personal monitoring of similar competing service offerings, because the resources thus spent would cut down any likely gains expected through the existing outsourcing agreement. On the other hand not monitoring offers at all means that our company would remain unable to seize opportunities for net savings through choosing less expensive service offers when it is advisable so to do.

It is clear therefore that the transformative development in such business contexts is the advent of automated assistance and choice making, if service offerings could be automatically monitored and their overall competitive position assessed, then the economic structures concerning such outsourcing decisions would be essentially changed. Just as in the telecommunication sector where long term commitments to monopolistic companies have been abandoned in favour of the now predominant call-by-call paradigm, here too our example company could be call-by-call enabled to choose the cheapest provider for every instance where a rental car, etc. is needed and keep the advantages of their existing contract plus the improved flexibility to choose a less expensive offer by a different provider, if possible. With such preferred "smart-sourcing" managed mix thus enabled through the co-existence of a basic agreement with a specific provider, and call-by-call service selection, then, the monitoring of less expensive offers would imply the invocation of the following simple heuristics to implement the appropriate trade-offs respectively in the following two cases:

- If the guaranteed amount of business is already fulfilled, always choose the least expensive offer for instant profits.
- If the guaranteed amount of business is not yet fulfilled:
  - Either choose cheaper offer and fulfil guarantee later (e.g. early in the guarantee period and if the gain is considerable);
  - Or choose outsourced offer to fulfil guarantee (e.g. towards end of guarantee period or if the outsider's rental rate offer is only mar-

marginal better anyway or in any case if the savings are outweighed by possible penalties applicable pursuant to non-fulfilment of the guaranteed business in any given period).

2.2. Semantic Web Services Monitoring

The basic interaction with Web services can be broadly divided into three stages [4]: the discovery of suitable services, the selection of one specific provider's service offering and the subsequent execution of the selected service (cf. figure 1). Starting with a service request a semantic description of the task is used for discovery of capable services. The candidate services are then queried based on the profile of the individual user and after an adequate quality assessment it should be possible to arrive at a choice of one service selected for subsequent execution. The task of automatically monitoring such services, and, service provision matchmaking thus comprises two phases: during the *discovery phase* all services that are basically capable of performing a certain task have to be accessed and queried for their best matching offers with respect to the specific task and during the *selection phase* these specific offers have to be matched against user preferences and the resulting offers have to be compared in terms of their quality.

[3] and [4] have in detail described algorithms to perform the above two phases with semantically enriched service descriptions in enhanced UDDI directories and queries expanded with information from given user profiles. An ontology-based approach is proposed providing a generalization hierarchy along which individual task requirements or query terms can be *relaxed* whenever no suitable services or offers can be found that exactly match the specified service requirement profile or a specific user's preferences.

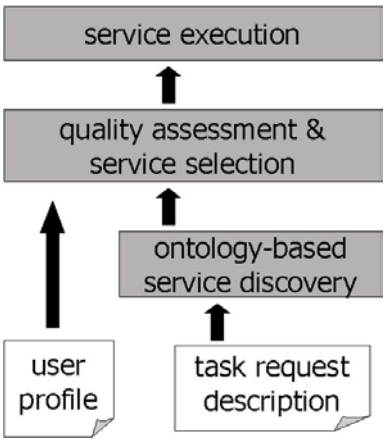

Figure 1: Stages of Web Service Interaction

These relaxation techniques play a crucial role in resolution of matchmaking choices and facilitating graceful closure behaviour in cooperative service provisioning i.e. in finding the best possible services/offers for execution. Just as with non-digital searches for suitable services, here too there can be no guarantee that available providers can always be found to satisfy a specific service request. For instance if a certain type of car is requested and our rental car company is just unable to make such a car available during the requested period, then the natural way is to relax the car-model-type constraint and propose to take a similar car type (for a thorough survey of such cooperative database retrieval techniques see [18]). This kind of relaxation can be performed automatically also in selecting Web services with certain capabilities. On the other hand each qualification component of a request also represents an objective of the user and thus relaxing several such qualifications can pose severe quality assurance problems. The multi-objective nature of the problem can lead to *fairness and closure hazards* i.e. incomparability between services/objects for which the request has been relaxed in different ways i.e. a differential embedded cost component may be involved across any units-of-quality thus relaxed; typical parameters like:

- a service's execution costs,
- the desirability of the offered objects / information,
- a service's trustworthiness or
- third party quality assessments / guarantees

have to be considered in the sense of Pareto-optimality, i.e. as maximal objects with respect to a (multi-dimensional) optimisation or compensation function [11].

Thus, although using the techniques in [3] and [4] enables users to have interesting alternative offers automatically monitored and still have the advantage of low administrative costs, the issue of how to assess and compare the quality of each service in each instance is still to be fully resolved. When a traditional outsourcing agreement is negotiated, the nature and quality of the service to be provided is fixed and thus a user familiar with a provider can rely on a certain quality guarantee in the context of a personal relationships based on implicit trust and confidence chains verified at some earlier stage through the collaborative service level agreement development lifecycle using various instruments in the tender and bidding process. In moving to a call-by-call paradigm with a vast variety of different providers for the adequate use of Web services we have to consider similar techniques and trust endorsement service structures and provisioning models to assure a service's usefulness.

3. Service Level Agreements and Web Service Quality

Using the ubiquity of the Internet with on-line services is, at least for many personal users, already part of today's life styles. A good example for first approaches to trust enhancement in service quality is practiced by managed service portals such as those provided by NTT DoCoMo's i-mode framework [19]. Here third party content providers offer services like news or sport results, routing/city maps or even car rental by registering with a portal managed by the network provider. Since the network provider also covers the billing and checks the services before they are provisioned via the portal, confidence in the individual content/service provider is enhanced. For use in WWW commerce, quality assurance structures such as those given in [10] or [17] have been proposed; these deal with similar trust endorsement challenges; usually involving intermediation by a "trusted third party" (like the network provider in the case of i-mode). Thus some means of endorsement, as in collaborative filtering, dispute resolution guarantees, licensing of service providers or insurance guarantees for potential customers, especially with respect to payments has been variously used to build confidence between customer and service provider. In a similar fashion [14] builds a model that allows users to choose between several implementations of Web services based on the services' reputation as e.g. given by evaluating feedback from previous users.

However, to achieve a solid foundation for a relationship between clients and service providers as in an outsourcing scenario we will need techniques that build not only confidence, but also a common understanding between outsourcer and service provider. In the following we will present a model for level of service agreements towards call-by-call outsourcing. Here we shall first define the term service level agreement and then show how the idea can be applied to real world scenarios.

Definition: A *Service Level Agreement* (SLA) is a mutual assurance instrument to support the delivery of services to be rendered meeting specific standards in terms of relevant attributes of vendor responsibilities guaranteed with respect to the cost of the service as agreed upon by the individual customer.

3.1. Formulating the Service Level Agreement

The formulation of such a mutual agreement in each instance requires the consideration of two tasks: Finding a service that is basically capable of performing the task or at least a service sufficiently similar to suit the task needed and then evaluating the availability of an adequate object that the service can offer. For instance though a service for car rental may be available, a specific type of car that is needed, may or may not be available independently of the service provider's capabilities. For example as in [3] for the discovery phase and in [4] for the selection phase; and also from a quality point of view we have to consider such mismatches differently each according to the nature of the mismatch. Let us first take a closer look at the discovery phase of applicable services. For the discovery part we have to find services that are generally able to help our problem.

Today modelling the concept of a service is mostly done within so-called UDDI directories [21] that provide a yellow pages style lookup mechanism giving a short semantic WSDL [5] description of a service together with its signature for invocation details. Advanced frameworks like WSMF [8] even supply some more details like necessary pre- and post-conditions for successful execution of a service, but finding the 'right' service to satisfy a specific need can still prove to be a difficult task. In [3] descriptive concept hierarchies stored as ontologies are proposed that generalize semantic descriptions of the services capabilities or attributes according to different user's conceptualisations. For example for a restaurant booking service the type of cuisine could be generalized along a regionally based taxonomy from Cantonese to other Chinese cuisines and further to related Asian cuisines and so on. Within the scope of this paper let us simply assume the existence of shared ontologies between the service provider and the service user. In other cases semantic Web techniques like DAML+OIL [6] or DAML-S [1] also allow us to deal with service requests where an ontology matching between users and providers has to be performed, e.g. [20], [7].

Applying the same relaxations with general Web services' capabilities needed for our service level agreements works exactly along the same lines as in the above example. Starting with the most specific required capability we can generalize it to related services that might still be able to serve our intended task. For example given our rental car use case scenario, if for whatever reason there should be no rental services available, we can generalize the specific task of 'car rental' along the concept of transportation to include e.g. flight booking, public transportation or taxi services or even general travel agents; unless there are specific hard constraints such as stipulated *class-exclusions* that must be observed for a particular user; for example "no air transport" for users with fear of flying. Of course the question of what amount of generalization we can sensibly allow, will strongly depend on the specific task and has to be decided upon in the context of the particular application domain. Anyhow by closely focussing on the problem this cooperative discovery behaviour is likely to lead to improvements upon today's restricted UDDI, even if a generalization strategy in some cases might not always prove useful.

Having found some services that are generally capable of performing our task we can proceed to the quality assessment and comparison amongst various competing service offerings. [4] describes an efficient method for selecting a service for subsequent execution based on comparing the diverse features that each service offers. It will retrieve the capable services' signatures and fill them with best matching values from a user's profile (like preferred cars) or generally applicable domain preferences (like cheapest price). Once this method is also integrated within our service selection platform, it will free the user from having to personally perform many tedious and time-consuming comparisons thus providing a supported call-by-call service provisioning environment like shown in figure 2. Let us first demonstrate the use of such an algorithm in a sample interaction scenario (section 3.2) and then turn to quality assessments to reduce the overhead for the comparisons amongst candidate service offerings (section 3.3).

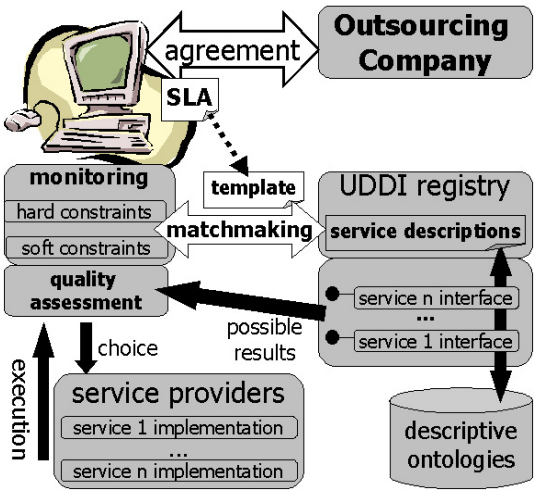

Figure 2: Call by Call Outsourcing Architecture

3.2. A Sample Use Case Query

Before we turn to the actual quality assessment, let us go through the phase of formulating the service level agreement with a sample query. Let us assume that a user has specified certain preferences about the car rental like pick-up location, rental period, etc. and some preferences about the car, like type or make. Usually a user's preferences list comprises both hard and soft constraints. So it is also necessary to distinguish between these two types of constraints; e.g. if a user has stated that they will need the car in the Bay Area then a relaxation of this location to other locations in California might not do, thus this is an example of a hard constraint. On the other hand a user might prefer a certain make of cars, if available, but will also accept another make; thus this implies that the car model/type can be processed as a soft constraint, i.e. relaxable if necessary. For simplicity we assume that the requester and provider share a common domain ontology or that anyhow the non-identical but equivalent terms (like 'rental_begin' or 'rental_start') can be matched automatically. Using the notation as described in [12] for a simple preference model featuring an 'I like A better than B' semantics in the form of strict partial orders and a preference-based declarative SQL-style query language, our user may formulate the basic request with hard and soft constraints as follows:

```
SELECT *
FROM UDDI_grounding('car rental service')
WHERE rental_location = 'San Francisco Bay Area'
    AND rental_begin = '10-05-03'
    AND rental_end = '10-09-03'
    AND car_non_smoking = 'true'
PREFERRING rental_pickup = 'San Francisco Airport'
          AND car_make = 'Volvo'
          AND min(rental_price)
```

This statement defines a certain location, car hire period and the fact that it has to be a non-smoking car as hard constraints and gives a preference for the pick-up location, a certain make and the price for the rental. If we take a closer look at these values and preferences we can see that for the discovery phase we would have to look for services whose grounding matches a semantic description like 'car rental service'. Assuming that an algorithm as described in [3] will come up with some adequate services then we have to focus on the differences between such candidate services and the subsequent selection based on a formulation of SLAs as reference pro forma contract and de facto service screening filter. The selection algorithm as described in [4] will now first filter out all services that do not supply all the necessary input parameters in their signature to allow the algorithm to ascertain if they can satisfy all the hard constraints. For the remaining services additional input parameters will be matched against the remaining preferences and suitable domain ontologies for relaxation can be chosen to see if any mismatches can be resolved (cf. figure 2).

If we look at the quality of the remaining unmatched preferences we are able to distinguish between different classes. For instance preferences like the pick-up location obviously have a geographic scope and can be relaxed along default ontologies consistent with common sense abstraction hierarchies. For example 'San Francisco Airport' is more specific than 'San Francisco' which in turn is more specific than 'Bay Area', and so on. It is generally more difficult to resolve the relaxation of preferences like the make of cars as may be desired by a user, because no assumption can be made re any car type taste and associated its relative preferences hierarchy as being a widely shared fact within the user community, hence no default relaxation taxonomies are applicable for such user

preferences. For this sort of personal taste, a personalised relaxation hierarchy has to have been stipulated by the user as part of their profile. We simply cannot decide for the user, whether for example, should no 'Volvo' be available, – they would rather have a 'BMW' or a 'Toyota' instead. But should a user explicitly express an interest in the car make, then it is possible also to transfer a short custom-made ontology together with the service request. The last preference about minimizing the price is not even an explicit preference, but results from the common understanding of the problem: Given two otherwise very similar or equal offers, it is a common sense and thus generalizeable assumption that customers always prefer the service with the lower price.

Having thus distinguished between different constraints and qualities of preferences we can formulate the first part of the service level agreement. Since the services are advertised as 'car rental service' we can assume a common understanding of this type of service provisioning and a widely shared expectation of what it means to enter into such a transaction and be in receipt of such a service including a generally accepted level of the operative standards of service, rates, roles and mutual responsibilities of the provider and the client. It is possible therefore to suggest that a tacit car hire service level agreement, in general is already in place on the basis of shared experience of established practice or social and business custom. And what is more, if a outsourcing agreement already has been made, i.e. the monitoring of new services is in competition with an existent agreement, we have a template of such an SLA by this outsourcing agreement already in place.

Additionally the client's hard constraints in each individual case have to be included in the agreement as basic facts which have to be provided if the situated personal service is to be properly delivered in each case. Please note that this may include some usual parameters like the car rental period that will differ from agreement to agreement, but will always be part of the agreement, plus some specific terms that a user may care about for some reason or other such as the non-smoking car in our case. Since not all people dislike smoking in cars, it would not be sensible to make such option a standard in every pro forma agreement that users have to consider in order to arrive at the *personalised situated SLA* that is to deliver the service required on each occasion. The explicit inclusion of such parameters in pro forma SLAs as standard structures thus eases their formulation, since it automatically supplies some standard business transaction parameters (that can reasonably be deemed to be part of common knowledge or shared social experience and expectation) and focuses on including what the user feels is necessary to be explicitly flagged in the SLA (e.g. "non-smoking" for people with allergies against cigarette smoke).

3.3. Assessing the Quality of Different Offers

After agreeing on the mandatory parts of the service level, i.e. standard service parameters and explicit hard constraints, we can focus on the rest of the service level agreement in terms of personal preferences. Since the hard constraints are mandatory their quality is basically assessed in a binary fashion: A constraint is fulfilled or not. Services with violated constraints can be removed from the candidate offers list. For instance considering our rental car scenario, if a service only offers rental cars for the area of Boston, they cannot be sensibly deployed to accomplish our task. In contrast, for the remaining soft constraints we have to focus on constraints that are more or less fulfilled, i.e. they show a certain grade of match without becoming inapplicable, if a constraint is violated within a certain degree. For example, if it should not be possible to pick up the car at the airport, depending on the user, picking it up in San Francisco or even somewhere in the Bay Area might still be acceptable, but accepting to do so somewhere in Northern California will usually be regarded as too broad a catchment area to be acceptable i.e. would have been *over-relaxed*. There are two ways to process such quality constraints:

- Users already within the request can provide a certain limit for the furthest possible relaxation for each parameter
- Users can view a ranked list of the best possible results and decide in each instance what is acceptable for them

As for providing a set of quality constraints already at the time of formulating the service request, again the values of an existing outsourcing agreement and its respective values can provide constraints that define maximum ranges or minimum standard demands according to a certain user's expectations.

Turning now to the issue of quality assessment, it is clear that that we have to decide between qualitative preferences and quantitative preferences see e.g. [9]. Quantitative preferences like the rental price in our example are defined over numerical domains and thus generally are easy to compare. Qualitative preferences like the location or car make are more complex to assess. [16] presents a framework to map qualitative domains onto quantifiable domains like for example our location could be easily mapped onto a distance between the San Francisco Airport and the exact pick-up location. Subsequently ordering and comparing these distances again is a simple task. So the only task left are the assessment of the qualitative preferences, which cannot be straightforwardly mapped onto numerical domains. Due to the openness of the preference model as strict partial orders even the proportional mapping onto a (normalized) numerical domain with the

top score for the best and the lowest score for the worst object in the preference hierarchy does not work, because incomparable values cannot always be mapped onto same score values as they would wrongly be transferred to comparables whereas in reality they are not so.

For this scenario [15] proposes a scheme of ordering different objects according to the preference hierarchies by counting the necessary relaxation steps from the top or the bottom of the hierarchy. Another way is to assume as distance the number of classes that an object dominates related to the number of classes that dominate it. Following these considerations the qualitative preferences can also be ordered and we can get to the point of selecting the best possible offers who will deliver our service level agreement. We have seen that within each attribute for which a preference has been specified, a comparison is possible, but we usually do not have an overall compensation function. Since we generally cannot decide for the user what is most important or how to compensate between attributes, i.e. if a convenient pick-up location can be considered more important than the rental price, we have to consider the attributes an equally important.

Following the notion of Pareto optimality [9] service offers can dominate other offers (i.e. considered to be better offers), only if they are better or at least of equal value than the other offers plus strictly better in respect of at least one quality criterion/preference element. Clearly since for our service level agreement we wish to arrive at the best possible offer for the requested set of user preferences, we can safely discard all dominated offers and propose the respective inclusion of soft constraints in all remaining offers as a possible closure for our service level agreement. The resulting *personalised situated SLA* would then finally contain

- all basic service standards as given characteristics,
- plus the necessary hard constraints of the user
- as well as all satisfiable user wishes.

Using cooperative techniques all user wishes in the SLA have been relaxed to the best matching objects that are available. Advanced preference-based systems as presented in [13] or [16] often also feature some added value in controlling the relaxation process by offering a possibility for the user to express a-priori in which order preferences should be relaxed and how far.

4. Summary and Outlook

In this paper we addressed the problem of assessing the quality of web service offers particularly in the context of online Call-by-Call outsourcing of services provision. Building on a scenario where a given task could possibly be accomplished by several services that may, however, offer different tradeoffs with respect to a user's needs and preferences, we considered ways of comparing such competing offers to inform initially a baseline description of available service characteristics to be developed into a mutual service level agreement as personalised for the user. From a business perspective such an agreement once accepted can be understood as a binding contract between the service requester and service provider; similar to an outsourcing relationship in the already established sense of SLA implementation that is practiced using the traditional outsourcing contracts. However unlike such traditional long-term outsourcing relationships this paradigm promises to pave the way towards a broader understanding and practice of the term outsourcing. It offers a complete outsourcing contract cycle, web-enabled and cost-efficient to be used for exploiting cost saving opportunities for suitable procurement of services rather than remaining captive to a long-term binding purchase agreement. Thus the client will be freed up to outsource afresh for such needs on a '*call-by-call*' manner. In each such instance the requester can easily decide whose offer should be taken, on the basis of the service level agreements. This is similar to the personal use of telephony services where by users can already practice call-by-call outsourcing routinely as they can make different calls through different telecommunication providers on the basis of their rate for a call. Essential for gaining this flexibility, however, is the monitoring and automatic comparison of different offers to reduce manual the administrative overhead attendant with call-by-call outsourcing.

After we identified the stages of interaction with the service and the necessary information in each stage, we presented a case study about how to formulate an adequate service level agreement in each instance. Our study started with the discovery of services generally capable of accomplishing the task and proposed semantic UDDI enhancements and service ontologies along the lines of [3] to build the 'common sense' foundations of a basic SLA. In the following selection phase we showed how to work all hard constraints that a user might be interested in, into the agreement utilising sophisticated algorithms to filter out all the services that would not meet the constraints. The selection step also used efficient techniques from [4] to match preferences and assess the quality for each offer leading to the derivation of the best possible service level agreement in each instance. Following the notion of Pareto optimality we then finally discarded all service offers that were dominated in all respects by other offers and thus returned a set of candidate agreements with the best possible fulfilment of all individual preferences as the final proposed service level agreement. This would become effective upon the issue of a commit in some form e.g. on issuing a purchase order.

Although this process will not always lead to an agreement, the cooperative determination of its content, which mirrors collaborative SLA development negotiations in traditional long-term outsourcing contracts, makes for improved flexibility and can thus be expected to meet the demands of modern e-commerce better than working through questionnaires with lots of possibly irrelevant issues as can occur with conventional agreements. Our future work will focus on the implementation of the ideas presented here, their economic evaluation and their integration into a suitable enterprise information system architecture taking the information of what an individual Web service can provide far beyond mere semantic descriptions in UDDI directories or simple WSDL statements. Given these mechanisms we believe that an improved usability of Web services to encompass also the area of B2C applications is now looming on the horizon.

Acknowledgements

We are grateful Matthias Wagner and the members of the e-Markets Research Group, University of Cambridge for their comments on our ideas. This research was funded within the Emmy-Noether program of the German Research Foundation (DFG).

References

[1] A. Ankolenkar, M. Burstein, J. Hobbs, et. al. DAML-S: Web Service Description for the Semantic Web. In *Proc. of the Int. Semantic Web Conf. (ISWC'02)*, Sardinia, Italy, LNCS 2342, Springer, 2002.

[2] A. Badii. Information Relationship Symmetry Management: a Critical Long-term Success Factor in Outsourcing ICT Systems and Services. In *Proc. of America's Conf. on Information Systems. (AMCIS'98)*, Baltimore, USA, 1998.

[3] W. Balke, M. Wagner. Cooperative Discovery for User-centered Web Service Provisioning. In *Proc. of the Int. Conf. on Web Services (ICWS'03)*, Las Vegas, USA, 2003

[4] W. Balke, M. Wagner. Towards Personalized Selection of Web Services. In *Proc. of the Int. World Wide Web Conf. (WWW 2003)*, Budapest, Hungary, 2003

[5] E. Christensen, F. Curbera, G. Meredith, S. Weerawarana. Web Services Description Language (WSDL) 1.1. http://www.w3.org/TR/2001/NOTE-wsdl-20010315, 2001.

[6] D. Connolly et al. DAML+OIL Reference Description. *W3C Note*, December 2001.

[7] A. Doan, J. Madhavan, P. Domingos, A. Halevy: Learning to map between ontologies on the semantic web. In *Proc. of the Int. World Wide Web Conf. (WWW 2002)*, Honolulu, Hawaii, USA, 2002

[8] D. Fensel and C. Bussler. The Web Service Modeling Framework WSMF. In *Journal of Electronic Commerce Research and Applications*. Elsevier, 2002.

[9] P. Fishburn. *Preference Structures and their Numerical Representations*. Theoretical Computer Science, 217:359-383, 1999

[10] M. Jakobsson, M. Yung. On Assurance Structures for WWW Commerce. In *Proc. of Int. Conf. on Financial Cryptography (FC'98)*, Springer LNCS 1465, Anguilla, British West Indies, 1998

[11] R. Keeney, H. Raiffa. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. Wiley & Sons, 1976

[12] W. Kießling, G.Köstler. Preference SQL: Design, Implementation, Experiences. In *Proc. of the Int. Conf. on Very Large Databases (VLDB'02)*, Hong Kong, China, 2002

[13] M. Lacroix, P.Lavency. Preferences: Putting more Knowledge into Queries. In *Proc. of the Int. Conf. on Very Large Databases (VLDB'87)*, Brighton, UK, 1987

[14] E. M. Maximilien, M.Singh. Conceptual Model of Web Service Reputation. In *SIGMOD Records* 31(4), 2002

[15] M. McGeachie, J. Doyle. Efficient Utility Functions for Ceteris Paribus Preferences. In *Proc. of Conf. on Artificial Intelligence and Conf. on Innovative Applications of Artificial Intelligence (AAAI/IAAI'02)*, Edmonton, Canada, 2002

[16] A. Motro. VAGUE: A User Interface to Relational Databases that Permits Vague Queries. In *ACM Transactions on Office Information Systems (TOIS)* 6(3), 1988

[17] G. Medvinsky, C. Lai, B. Neuman. Endorsements, Licensing, and Insurance for Distributed System Services. In *Proc. of the ACM Conf. on Computer and Communications Security*, Fairfax, USA, 1994

[18] J. Minker. An Overview of Cooperative Answering in Databases. In *Proc. of the Int. Conf. on Flexible Query Answering Systems (FQAS)*, Roskilde, Denmark, LNCS 1495, Springer, 1998

[19] NTT DoCoMo Homepage http://www.nttdocomo.com/home.html, 2003

[20] M. Paolucci, T. Kawamura, T. Payne, K. Sycara. Importing the Semantic Web in UDDI. In *Proc. of the Int. Workshop on Web Services, e-Business and the Semantic Web (WES'02)*, Toronto, Canada, 2002

[21] UDDI. The UDDI Technical White Paper. http://www.uddi.org.

[22] M. Wagner, W.-T. Balke, R. Hirschfeld, W. Kellerer. A Roadmap to Advanced Personalization of Mobile Services. In *Proc. of the Int. Conf. DOA/ODBASE/ CoopIS (Industry Program), Irvine,* USA, 2002